



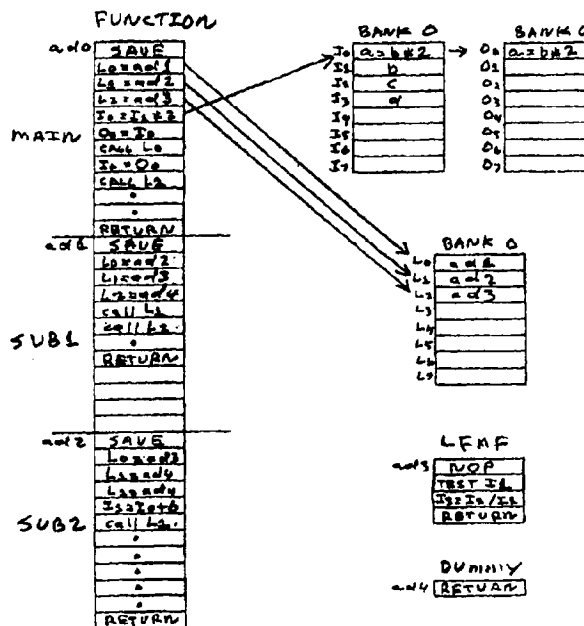
INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(15) International Patent Classification ⁶ : G06F 15/18		A2	(11) International Publication Number: WO 98/02825
			(43) International Publication Date: 22 January 1998 (22.01.98)
(21) International Application Number: PCT/US97/11905		(81) Designated States: AM, AT, AU, BB, BG, BR, BY, CA, CH, CN, CZ, DE, DK, EE, ES, FI, GB, GE, HU, IS, JP, KE, KG, KP, KR, KZ, LK, LR, LT, LU, LV, MD, MG, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, TJ, TM, TT, UA, UG, UZ, VN, ARIPO patent (GH, KE, LS, MW, SD, SZ, UG, ZW), European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG).	
(22) International Filing Date: 10 July 1997 (10.07.97)			
(30) Priority Data:			
08/679,555	12 July 1996 (12.07.96)	US	
08/682,859	12 July 1996 (12.07.96)	US	
08/674,337	12 July 1996 (12.07.96)	US	
(71)(72) Applicant and Inventor: FRANCONE, Frank, D. [US/US]; 360 Grand Avenue #355, Oakland, CA 94610 (US).		Published Without international search report and to be republished upon receipt of that report.	
(72) Inventors: NORDIN, Peter ; Dr. Westringssgain 21 L, S-413 23 Göteborg (SE). BANZHAF, Wolfgang ; Ackerstrasse 16, D-44575 Castrop-Rauxel (DE).			
(74) Agent: ALEXANDER, David, G. ; Bronson, Bronson & McKinnon, Suite 2400, 444 South Flower Street, Los Angeles, CA 90071-2925 (US).			

(54) Title: COMPUTER IMPLEMENTED MACHINE LEARNING METHOD AND SYSTEM

(57) Abstract

One or more machine code entities (50) such as programs or functions are created which represent solutions to a problem and are directly executable by a computer system (10). The programs (50) are created and altered by a program (32) in a higher level language such as "C" which is not directly executable, but requires translation into executable machine code through compilation, interpretation, translation, etc. The entities (50) are initially created as an integer array (32b) that can be altered by the program (32) as data, and are executed by the program (32) by recasting a pointer to the array as a function type. The entities (50) are evaluated by executing them with training data as inputs, and calculating fitnesses based on a predetermined criterion. The entities (50) are then altered based on their fitnesses using a genetic machine learning algorithm by recasting the pointer to the array as a data (e.g. integer) type. This process is iteratively repeated until an end criterion is reached. The entities (50) evolve in such a manner as to improve their fitness, and one entity is ultimately produced which represents an optimal solution to the problem. Each entity (50) includes a plurality of directly executable machine code instructions (52a, 52b, 52c), a header (50a), a footer (50c), and a return instruction (50d). The instructions (50) can include branch instructions which enable subroutines, leaf functions, external function calls, recursion, and loops. The system (10) can be implemented on an integrated circuit chip (90), with the entities (50) stored in high speed memory (96) in a central processing unit (92). The system (10) can be used to control an autonomous agent such as a robot.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav	TM	Turkmenistan
BF	Burkina Faso	GR	Greece		Republic of Macedonia	TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's	NZ	New Zealand		
CM	Cameroon		Republic of Korea	PL	Poland		
CN	China	KR	Republic of Korea	PT	Portugal		
CU	Cuba	KZ	Kazakhstan	RO	Romania		
CZ	Czech Republic	LC	Saint Lucia	RU	Russian Federation		
DE	Germany	LI	Liechtenstein	SD	Sudan		
DK	Denmark	LK	Sri Lanka	SE	Sweden		
EE	Estonia	LR	Liberia	SG	Singapore		

5 COMPUTER IMPLEMENTED MACHINE LEARNING METHOD AND SYSTEMTechnical Field

The present invention generally relates to the art of computerized computation systems for performing repeated computations on data that is not known until a computer user is running the system ("run-time", and "run-time data") and more specifically to a system of creating, initializing, storing, altering and executing both
10 the run-time data and the computer code necessary to execute the repeated computations (the "related code") in native machine code operating on a register machine. The present invention relates also to the art of computerized learning systems (which are usually characterized by the need to perform repeated computations on run-time data) and more specifically to a register machine learning method in which the information and/or the computer program(s) that constitute solutions to a problem are created, initialized, stored, altered and executed in native
15 machine code by using a higher level programming language to produce an optimized solution through the direct application of learning algorithms to the information stored in the native machine code.

Background Art

Machine learning systems have been proposed in the art for the solution of problems such as classification, prediction of time-series data, symbolic regression, optimal control, etc. Examples of various
20 machine learning systems are neural networks, fuzzy networks, genetic algorithms (including genetic programming and classifier systems), Evolutionary Strategies, Evolutionary Programming, ADAPTIVE program induction, cellular automata, Box Jenkins optimization, ARMA optimization and many others. Rather than applying a direct computational approach, these systems create one or more proposed solutions in the form of data and computer program entities, and iteratively alter the data and/or entities for the purpose of finding an optimal solution to a
25 problem. One such approach is described in, for example, U.S. Patent No. 4,935,877, entitled "NON-LINEAR GENETIC ALGORITHMS FOR SOLVING PROBLEMS", issued June 19, 1990 to John Koza.

The set of practically solvable problems is highly related to the efficiency of the algorithm and implementation. It is therefore important to minimize the overhead involved in executing any machine learning system.

30 Machine learning systems create learned solutions to problems. These solutions have two elements: (1) Elements that hold the information that is learned during the execution of the machine learning system (the "Learned Elements"); and (2) Elements that are necessary to convert the Learned Elements into meaningful action or output by the computer (the "Conversion Elements"). Existing machine learning approaches other than the present invention can be classified into two categories in terms of how they store and execute both the Learned
35 Elements and the Conversion Elements. Those two categories are compiler based systems and interpreted systems.

An interpreted system is written in a high level language such as LISP and both the Learned and Conversion Elements are held in data-like structures such as LISP lists, which are converted ("interpreted") into native machine code at run-time by an interpreter. The interpreter, itself, also contains Conversion Elements for such interpreted systems. So for example, U.S. Patent No. 4,935,877 uses, as its Learning Elements various high

level LISP expressions, customized for the problem at hand, to represent, symbolically, a "computer program." That is, a high level "program" structure symbolized by the LISP List, is itself the subject of learning in that system. In this system, the Learned Elements are represented as a hierarchical tree structure. This solution gives good flexibility and the ability to customize the language depending on the constraints of the problem at hand.

5 The principal disadvantage of this interpreting approach to machine learning is that the Learned Elements and many of the Conversion Elements are stored in high level, symbolic data-like structures. Computers can operate only by executing native machine code. Thus, interpreting machine learning systems learn by modifying high level symbolic representations (the Learned Elements) that are, ultimately, converted into machine code by the interpreter at run-time. The need to convert (interpret) the Learned Elements and some of the
10 Conversion Elements into native machine code at run-time before any useful action or output may be had from a computer is very time consuming and involves a large amount of overhead in machine resources such as CPU time, RAM memory, and disk space. In effect, all of the Learned Elements and the Conversion Elements in the LISP List are treated as run-time data that must be accessed and converted to machine code before any useful action or output may be had. Simply put, interpreted systems are slow and use a lot of a computer system's
15 resources.

Other machine learning systems are written in high level compiled computer languages such as C, C++, Pascal and so forth. This is the "compiler based" approach to machine learning. Large amounts of the Conversion Elements in such systems are compiled before run-time into native machine code in such a compiler based approach. Because those compiled Conversion Elements are already in native machine code at run-time, there is
20 no need to interpret these Conversion Elements at run-time. Using a compiler based approach instead of an interpreted approach usually results in a substantial speed-up of the execution of the machine learning system, often increasing the speed of learning by a factor of ten times. Examples of such compiler based systems are GPC by Walter Tackett (genetic programming), Lil-GP (genetic programming), and EvoC (evolutionary strategies).

The compiler based approach to machine learning, while faster than the interpreted approach, must
25 still access run-time data that is stored in data structures held in RAM memory or in some other form of memory such as hard disk. The reason that run-time data structures must be accessed in compiler based machine learning systems (or any machine learning system other than the present invention) is that the process of learning involves initializing and then altering the Learned Elements at run-time.

For example, the weights in a neural network are Learned Elements for neural network applications.
30 Compiler based neural network systems hold those weights in data structures such as arrays or linked lists. Similarly, compiler based genetic programming systems store symbolic representations of program structures (the Learned Elements in genetic programming) in RAM data structures such as arrays, linked lists or parse trees. In all compiler based machine learning systems, the already compiled Conversion Elements must repeatedly access the Learned Elements (weights or symbolic representations of program structures) from the data structures in RAM
35 memory in order to execute and evaluate the Learned Elements and to modify the Learned Elements according to the learning algorithm that is being used to modify the Learned Elements during learning. Such repeated access is necessary before any meaningful output or action may be had from a computer based on the Learned Elements.

Such repeated accesses to RAM data structures is time consuming and uses extensive amounts of RAM to hold the Learned Elements. More generally, computing systems that perform repeated calculations on run-time data may also be categorized as compiler based systems or interpreted systems. They store access, alter and execute run-time data in a manner similar to the storage, access, alteration and execution of the Learned Elements in the systems described above and are subject to the same limitations of slow execution and system overhead as the systems described above. By the phrase, "repeated calculations (or computations) on (or of) run-time data," this application means the execution of one or more instructions that must access one or more elements of run-time data (from data storage such as RAM or hard disk) more than once on the same values of the run-time data.

DISCLOSURE OF INVENTION

10 The present invention utilizes the lowest level binary machine code as the "entities" or individuals or solutions. Every individual is a piece of machine code that is called and manipulated by the genetic operators.

There is no intermediate language or interpreting part of the program. The machine code program segments are invoked with a standard C function call. Legal and valid C-functions are put together, at run time, directly in memory by the genetic algorithm. The present system can therefore be considered as a "compiling" machine learning implementation.

The present system generates binary code directly from an example set, and there are no interpreting steps. The invention uses the real machine instead of a virtual machine, and any loss in flexibility will be well compensated for by increased efficiency.

More specifically, one or more machine code entities such as functions are created which represent
20 (1) solutions to a problem or (2) code that will perform repeated calculations on "run-time data" that is encapsulated into the machine code. These entities are directly executable by a computer. The programs are created and altered by a program in a higher level language such as "C" which is not directly executable, but requires translation into executable machine code through compilation, interpretation, translation, etc.

The entities are initially created as an integer array that can be altered by the program as data, and
25 are executed by the program by recasting a pointer to the array as a function type. The entities are evaluated by executing them with training data (as defined elsewhere) as inputs, and calculating "fitnesses" based on a predetermined criterion or by recovering the output as the result of one of the repeated calculations on run-time data.

After one or more "executions," the entities are then altered by recasting the pointer to the array as
30 a data (e.g. integer) type. Or the original data pointer to the array may have been typecast earlier to a function pointer in a way that did not permanently change the type of the pointer. In that case, the original data pointer is used in its original form. This process is iteratively repeated until an end criterion is reached.

In the case of machine learning, the entities change in such a manner as to improve their fitness, and one entity is ultimately produced which represents an optimal solution to the problem. In the case of repeated
35 calculations on run-time data, the entity permits very rapid calculations to be performed generating usable output from the run-time data.

Each entity includes a plurality of directly executable machine code instructions, a header, a footer, and a return instruction. The alteration process is controlled such that only valid instructions are produced. The headers, footers and return instructions are protected from alteration.

The system can be implemented on an integrated circuit chip, with the entities stored in high speed
5 memory in a central processing unit.

The present invention overcomes the drawbacks of the prior art by eliminating all compiling, interpreting or other steps that are required to convert a high level programming language instruction such as a LISP S-expression into machine code prior to execution.

Experiments have demonstrated that the present approach can speed up the execution of a machine
10 learning or a repeated calculation system by 1,000 times or more as compared to systems which provide potential solutions in the form of high level "program" expressions in the interpreted approach. The speedup is in excess of 60 times over the compiler based approach. This makes possible the practical solutions to problems which could not heretofore be solved due to excessive computation time. For example, a solution to a difficult problem can be produced by the present system in hours, whereas a comparable solution might take years using
15 conventional techniques.

These and other features and advantages of the present invention will be apparent to those skilled in the art from the following detailed description, taken together with the accompanying drawings, in which like reference numerals refer to like parts.

BRIEF DESCRIPTION OF DRAWINGS

20 FIG. 1 is a block diagram illustrating a machine learning system according to the present invention as implemented on a general purpose computer;

FIG. 2 is a block diagram illustrating a computer program which implements the functionality of the present invention as stored in a random access memory of the computer of FIG. 1;

FIG. 3 is a block diagram illustrating banks of registers in the computer used in the preferred
25 implementation of the invention;

FIG. 4 is a flowchart illustrating a method of machine learning according to the invention;

FIG. 5 is a diagram illustrating types of machine code instructions which are used by the preferred implementation of the invention;

FIG. 6 is a detailed flowchart illustrating a specific embodiment of the preferred implementation of
30 the invention;

FIG. 7 is a diagram illustrating an array of functions that are used by the invention;

FIG. 8 is a diagram illustrating an alternative function that can be used by the invention;

FIG. 9 is a diagram illustrating the present machine learning system as implemented on a specially designed integrated circuit chip;

35 FIGs. 10a and 10b are diagrams illustrating a genetic uniform crossover operation;

FIGs. 11a and 11b are diagrams illustrating a genetic 2-point "swap" crossover operation;

- FIGs. 12a and 12b are diagrams illustrating a genetic operator crossover operation;
FIGs. 13a and 13b are diagrams illustrating a genetic operand crossover operation;
FIG. 14 is a diagram illustrating a genetic operator mutation operation;
FIG. 15 is a diagram illustrating a genetic operand mutation operation;
5 FIGs. 16a to 16d in combination constitute a flowchart illustrating a generic implementation of the invention to machine learning;
FIGs. 17a to 17d in combination constitute a flowchart illustrating a generic implementation of the invention to repetitive computations based on run time data;
FIGs. 18a to 18h in combination constitute a flowchart illustrating an machine learning
10 implementation of the system which induces machine code functions where the Learned Elements are machine code instruction structure and machine code instruction contents. This is one implementation of what is referred to loosely as Compiling Genetic Programming Systems;
FIGs. 19 to 21 are diagrams illustrating the use of functions and registers in accordance with the invention;
15 FIGs. 22a to 22k are flowcharts illustrating a detailed implementation of the invention;
FIG. 23 is a diagram illustrating non-memory control of an autonomous agent such as a robot according to the present invention;
FIG. 24 is a flowchart illustrating the non-memory control process of FIG. 23;
FIGs. 25 and 26 are diagrams illustrating control of an autonomous agent using a learning unit and
20 a memory; and
FIGs. 27 and 28 are flowcharts illustrating the control process of FIGs. 25 and 26.

MODE(S) FOR CARRYING OUT THE INVENTION

The present method must be implemented using a computer due to the immense number of complex computations that must be made. The computer can be a general purpose computer, with the functionality of the
25 method being implemented by software. Alternatively, as will be described below, part or all of the functionality of the system can be implemented on a specially designed integrated circuit chip.

Introduction

The present invention utilizes the lowest level native machine code with no immediate step of compilation or interpretation to store, access, initialize, create, alter and execute run-time data and related code
30 where repeated computations (embodied in the related code) must be performed using that run-time data. In the context of machine learning, the run-time data (the Learned Elements) and the computations that are to be performed on the run time data (the Conversion Elements) are created, initialized, stored, altered and executed directly in native machine code with no intermediate step of compilation or interpretation. Thus, in the context of machine learning, the present invention stores both Learned Elements and the Conversion Elements in native
35 machine code. All or many of the Learned Elements and all of the Conversion elements are created, initialized,

stored, altered and executed directly in the native machine code with no intermediate step of compilation or interpretation.

The present invention is not limited to any particular machine learning paradigm such as, and by way of example only, genetic programming, genetic algorithms, simulated annealing or neural networks. Rather, it is a method of creating, initializing, storing, altering and executing all or part of the Learned Elements and the Conversion Elements in any machine learning paradigm directly in native machine code.

For example, when the present invention is applied to evolving native machine code structure and contents, it creates, initializes, stores, alters and executes the programs structures that are to be evolved (the Learned Elements) directly in native machine code. The approach of the present invention is completely unlike the compiler based and interpreting approaches to evolving "computer programs," which only create, store, alter and (with the aid of a compiler or an interpreter) execute high level symbolic representations of high level "program structures" in a high level programming language that, ultimately, represent and are converted into actual machine code by an interpreter or that are executed by compiled code making repeated accesses to the run-time Learned Elements (the representations of high level "program structures") in RAM.

As a further example, when the present invention is applied to neural networks, it creates, initializes, stores, alters and executes the weights of the network (the Learned Elements) and the code that executes a network characterized by those weights (the Conversion Elements) directly in native machine code. This is completely unlike the compiler based and interpreting approaches to neural networks, which create, initialize, store, alter and (with the aid of a compiler or an interpreter) execute high level symbolic representations of weights and network structures that, ultimately, represent and are converted by an interpreter into the actual machine code that executes the neural network or that are executed by compiled code making repeated accesses to the run-time Learned Elements in RAM.

The present invention utilizes functions created in the lowest level machine code as the "learning entity" for machine learning or as the "computation entity" for repeated computations based on run-time data. Every such learning entity or computation entity (also referred to as an "entity" or an "individual" or a "solution") in the present invention is a complete, executable native machine code function that can be executed from a higher level language such as C or assembler by a simple standard function call ("Native Functions"). All of the run-time data and related code, the Learned Elements and the Conversion Elements are assembled by the present invention at run time into Native Functions, which constitute the learning entity or the computation entity, and the learning entity or computation entity are then stored, altered and executed directly as Native Functions. When the run-time data or the Learned Elements change, the learning entities or the computation entities must also be changed because the previous run-time data, including the previous Learned Elements, are already included in the Native Function learning entity. The present invention is, in part, a method of making such changes directly in the learning and computation entities. In a sense, the present system acts as an on-the-fly compiler of both real time data and related code (for computation entities), and of Learned Elements, and Conversion Elements (for learning entities). The present system can therefore be loosely considered as a "compiling" machine learning implementation. In the specific case of evolving machine code structures and contents as the Learned Elements,

the present methodology can be loosely described as a "Compiling Genetic Programming System (CGPS)".

The present system generates low level machine code directly from an example set, and there are no interpreting or compiling steps. The invention uses the real machine instead of a virtual machine, and any loss in flexibility will be well compensated for by increased efficiency.

5 More specifically, one or more individuals (or "entities" or "solutions") are created which represent solutions to a problem and are directly executable by a computer as Native Functions. The programs are created, initialized, stored, altered and executed by a program in a higher level language such as "C" which is not directly executable, but requires translation into executable machine code through compilation, interpretation, translation, etc.

10 The entities are initially created as an integer array that can be created, initialized, stored and altered by the higher level program as data. Although the present implementation of the system uses arrays of integers, the system could be implemented as arrays of any other data type and such implementations are within the scope and spirit of the present invention.

To understand how the system chooses what integers are included in this array, one must understand
15 what a low level machine instruction looks like. The low level instructions on all existing CPU's are groups of binary bits. On the Sun system for which most of the development of the present system has taken place, all instructions are 32 bits long. In low level machine code, each bit has a particular meaning to the CPU. In the present system, the integers chosen for the integer arrays that constitute the entities are integers that, when viewed in binary format, correspond to valid low level machine code instructions for the CPU used in the computer system
20 to which the present invention is being applied. Although the present system is implemented on Sun and Intel systems, it could be implemented on any CPU and many other chips. Such alternative implementations are within the scope and spirit of the present invention.

In the present implementation of the invention in C, the entities are executed by typecasting a pointer to the array of integers that constitute the entity to be a pointer to a function. The function pointed to by the
25 entity is then executed by calling the function using the function pointer like a regular C function. The important point here is that the entities are viewed, alternatively, as data and as functions at different times in the system. At the time the entities are created, initialized, stored and altered, the entities are viewed by the system as data-an array of integers. At the time the entities are executed, they are viewed as Native Functions. So if the present system was implemented in assembler, the entities would be created as integers (as data) and they would be
30 executed as Native Functions. No typecast would be required in assembler but such an implementation is within the scope and spirit of the present invention.

When the entities are executed, they use training data (defined below) as inputs. The results of the execution are converted into some measure of how well the problem is solved by the entity, said measure being determined by the machine learning paradigm or algorithm being used. Said measure will be referred to herein
35 as "fitness" regardless of the machine learning paradigm or algorithm being used. Many machine learning algorithms take the measure of fitness and feed it back into the algorithm. For example, evolutionary based learning algorithms use fitness as the basis for selecting entities in the population for genetic operations. On the

other hand simulated annealing uses fitness, inter alia, to determine whether to accept or reject an entity as being the basis for further learning. After entities are initialized, most machine learning algorithms, go through repeated iterations of alteration of entities and execution of entities until some termination criterion is reached.

In the particular implementation of the present invention, the entities are initially created and
5 initialized as an integer array that can be altered by the program as data and are executed by the program by recasting a pointer to the array as a pointer to a function type. The entities are evaluated by executing them with training data as inputs, and calculating fitnesses based on a predetermined criterion applied to the output of an entity after execution.

The entities are then altered based on their fitnesses using a machine learning algorithm by recasting
10 the pointer to the array as a data (e.g. integer) type or by using the original data pointer. This process is iteratively repeated until an end criterion is reached.

The entities evolve in such a manner as to improve their fitness, and one entity is ultimately produced which represents the best solution to the problem.

Each entity includes a plurality of directly executable machine code instructions, a header, a footer,
15 and a return instruction. The alteration process is controlled such that only valid instructions are produced. The headers, footers and return instructions are protected from alteration.

The system can be implemented on an integrated circuit chip, with the entities stored in high speed memory or in special on chip registers in a central processing unit and with the creation initialization, storage, alteration and execution operators stored in chip based microcode or on ROM.

20 The present invention overcomes the drawbacks of the prior art by eliminating all compiling, interpreting or other steps that are required to convert a high level programming language instruction such as a LISP S-expression into machine code prior to execution. It overcomes the need of compiler based systems to access, repeatedly, the Learned Elements as run-time data, which greatly slows down the system. It also represents the problem to the computer in the native language of the computer, as opposed to high level languages that were
25 designed for human, not for machine, understanding. It acts directly on the native units of the computer, the CPU registers, rather than through the medium of interpreters or compilers. It is believed that by using an intermediate process such as interpretation or compilation of high level codes, previous systems may actually interfere with the ability of computers to evolve good solutions to problems.

Experiments have demonstrated that the present approach can speed up the execution of a machine
30 learning system by 1,000 times or more as compared to systems which provide potential solutions in the form of high symbolic level program expressions and by 60 times or more as compared to compiler based systems. This speed advantage makes possible practical solutions to problems which could not heretofore be solved due to excessive computation time. For example, a solution to a difficult problem can be produced by the present system in hours, whereas a comparable solution might take years using conventional techniques.

35 These and other features and advantages of the present invention will be apparent to those skilled in the art from the following detailed description, taken together with the accompanying drawings, in which like reference numerals refer to like parts.

Detailed Implementation

FIG. 1 illustrates a computer system 10 for implementing a machine learning method according to the present invention. The system 10 includes a Random Access Memory (RAM) 12 in which the software program for implementing the functionality of the invention is stored, and a processor 14 for executing the
5 program code. The system 10 further includes a visual display unit or monitor 16 for providing a visual display of relevant information, a read-only memory (ROM) 18 for storing firmware, and an input-output (I/O) unit 10 for connection to a printer, modem, etc.

The system 10 further includes a mass data storage 20 which can be any combination of suitable elements such as a fixed (hard) magnetic disk drive, a removable (floppy) disk drive, an optical (CD-ROM) drive,
10 etc. Especially large programs which implement the invention may be stored in the storage 20, and blocks of the programs loaded into the RAM 12 for execution as required.

User access to the system 10 is provided using an input device 22 such as alphanumeric keyboard 114 and/or a pointing device such as a mouse (not shown). The elements of the system 10 are interconnected by a bus system 24.

15 The system 10 is preferably based on the SUN SPARC architecture due to its register structure. The invention is also preferably implemented using the "C" programming language due to its freedom in the use of data types and the ability to cast between different types, especially pointers to arrays and pointers to functions and a standard SUN operating system compiler, although the invention is not limited to any particular configuration.

20 The invention has been practiced using the above configuration in the SPARC environment which generally a stable architecture for low level manipulation and program patching. The particular platform used was the SUN SPARCSTATION 1+.

The RAM 12 is illustrated in FIG. 2 as storing an operating system 30 such as UNIX or one of its variants, and a program 32 written in the "C" language which provides the functionality of the invention. The
25 program 32 includes high level machine language instructions 32a, and one or more machine code array 32b.

The high level instructions 32a are not directly executable by the processor 14, but must be converted into directly executable machine code by a compiler. The array 32b, however, is provided in native (lowest level) machine code including binary instructions that are directly recognized and executed by the processor 14.

The performance of the present invention is greatly enhanced by implementation thereof using a
30 processor having banks of registers. This architecture is provided by the SPARC system as illustrated in FIG. 3.

The components of the SPARC processor 14 which are most relevant to understanding the present invention include an Arithmetic Logic Unit (ALU) 40 for performing actual numeric and logical computations, a program counter 42 for storing an address of a next instruction to be executed, and a control logic 44 which controls and coordinates the operations of the processor 14 including memory access operations.

35 The processor 14 further includes banks of registers which are collectively designated as 46. The processor 14 also includes a number of additional registers which are not necessary for understanding the principles of the present invention, and will not be described. Although the invention is presently implemented using integer

arithmetic, it may be applied to Floating Point arithmetic easily and such an application is within the scope and spirit of the invention.

More specifically, the registers 46 include eight register banks (only one is shown in the drawing), each of which has 8 output registers O_0 to O_7 , 8 input registers I_0 to I_7 , 8 local registers L_0 to L_7 . There are, also, 8 global registers G_0 to G_7 , accessible from any bank. The implementation of the invention using this architecture will be described below.

FIG. 4 is a flowchart illustrating the general machine learning method of the present invention which is implemented by the high level instructions 32a in the C programming language. Figures 22a to 22k are more detailed flowcharts illustrating the present invention. The first step is to define a problem to be solved, a fitness criterion for evaluating the quality of individual solutions, and an end criterion. The end criterion can be a predetermined number of iterations for performing the process (e.g. 1,000), achievement of a predetermined fitness level by one or more solutions, a change in total or individual fitness which is smaller than a predetermined value, or any other suitable criterion.

Next, the input, output, calculation, and state registers must be identified.

An input register is a register that is initialized before each calculation with data. In the present implementation of the invention, the input registers may be any one of the registers referred to as I0 through I5 in the Sun architecture. See Figure 3, number 46. In the present implementation of the invention, it is, therefore possible to have up to six input registers. However, it would be easy to extend the system to have many more input registers and such an extension is within the scope and spirit of the present invention. In any event, the designer must pick how many of the eight input registers to use.

There are many ways to initialize these registers but the preferred implementation is to include the input variables as parameters when the individual is called as a Native Function. For example, if native_function_array is an array of integers that constitute valid machine code instructions that amount to a Native Function and that amount to an entity in this invention, then the following type definition and function call execute the native_function_array individual and pass two input variables to it.

```
Typedef unsigned int(* function_ptr) ();
{the type definition}
a=((function_ptr) native_function_array)(2,3);
{calls the function and passes two inputs (2 and 3), which are placed in I0 and I1.}
```

The items in curly brackets are descriptive and not part of the code. In the present implementation of this invention, the two input variables are put into I0 and I1. Accordingly, if the designer decides to have two input registers, they must be I0 and I1 in the present implementation. Other methods of initializing registers such as putting instructions in the header of the entity (see Figure 5, number 50) that retrieve data from RAM provide more flexibility as to what registers are input registers and more input registers and they are within the scope and spirit of the present invention.

This process of picking and initializing input registers is referred to in Figures 16b, c, and d; 17b, c, and d; 18b, d, and f; and 22b, e, and g.

An output register is read after the individual has finished executing on a given set of inputs. In other words, an output register is the output of the individual on a given set of inputs. The designer must designate one or more registers to be the output register(s) and the number depends on the problem to be solved. In the present implementation, the value in I0 at the end of the execution of the individual is automatically set as the single output register and as the output of the individual. The above code automatically puts the value in I0 into the variable "a." It is possible to designate multiple output registers and to preserve the values therein when an individual has been executed. One such method would be to include instructions in the footer of the individual (see Figure 5, number 50) that move the value(s) in the designated output registers in the individual that has just executed to storage registers or to locations in RAM memory. Such implementations are within the scope and spirit of the present invention.

This process of picking and recovering the values in output registers is referred to in Figures 16b, c, and d; 17b, c, and d; 18b, d, and f; and 22b, e, and g.

A calculation register is a register that is neither an input nor an output register but that is initialized to a given value every time the individual is executed. Zero or one is the preferred initialization value. A calculation register may be initialized in the headers of the individuals. (See Figure 5, number 50). In any event the designer must decide how many, if any, calculation registers he or she wishes to use and then designate specific registers to be calculation registers. So, for example, if there were two input registers, I0 and I1 and if I0 were the output register, then I2 through I5, and L0 through L7, among others would be available slots for a calculation register. So if the designer desires one calculation register, he could pick I2 as the calculation register.

This process of picking and initializing calculation registers is referred to in Figures 16b, c, and d; 17b, c, and d; 18b, d, and f; and 22b, e and g.

A state register is a register that is neither an input nor an output register but that is not initialized every time the individual is executed. That is, values from some previous execution of the individual are retained in a state register. The precise retention policy is a matter of decision for the designer. At the beginning of the execution of the individual, the state register has to be initialized to the value it held at the end of a previous execution of the individual. If that value was saved in a storage register or in RAM, then the header of the individual (See Figure 5, number 50) must have an instruction moving the value from the storage register or RAM to the state register or the previous state could be passed to the individual as a parameter when the individual is called as a Native Function. In this case, the footer of the individual may have an instruction saving the value of the state register at the end of execution to the storage register or to RAM. There are many available ways to implement state registers and to save and initialize their values and all of them are within the scope and spirit of this invention. In any event the designer must decide how many, if any, state registers he or she wishes to use and then designate specific registers to be state registers. So, for example, if there were two input registers, I0 and I1, and one calculation register, I2, and if I0 were the output register, then I3 through I5, and L0 through L7, among others would be available for a state register. So if the designer desires one state register, he or she could

pick 13 as the state register.

This process of picking and initializing state registers is referred to in Figures 16b, c, and d; 17b, c, and d; 18b, d, and f; and 22b, e and g.

Once all registers have been designated, the designer has established the "Register Set." In the
5 example above, the Register Set would be as follows

input	10, 11
output	10
calculation	12
state	13

10 Once the Register Set has been picked, it is essential that all individuals be initialized with instructions that contain only references to valid members of the register set. See Figures 16c, 17c, 18d, and 22e. It is also essential that all changes to the individuals, including those made with genetic operators, modify the individual in such a way as to make sure that all register references in the modified instructions are to valid members of the Register Set.

15 Next, a solution or population of solutions is defined as an array of machine code entities. These entities are preferably "C" language functions which each include at least one directly executable (machine language) instruction. An instruction is directly executable by being constituted by the lowest level (native) binary numbers that are directly recognized by the control logic 44 and ALU 40 as valid instructions. These instructions do not require any compilation, interpretation, etc. to be executable.

20 The invention will be described further using the exemplary case in which the entities are "C" language functions. However, it will be understood that the invention is not so limited, and that the entities can be constituted by any data structures including machine code instructions that can alternatively be manipulated as data and executed by a higher level program.

After the function or array of functions has been defined, they are initialized by inserting valid
25 machine code instructions in the appropriate locations therein. As illustrated in FIG. 5, a function 50 includes a header 50a, an instruction body 50b, a footer 50c, and a return instruction 50d.

The header 50a deals with administration which is required when entering a function. This normally involves manipulation of the registers 46, including passing arguments for the function to and from the various registers 46, saving the values in the registers that exist when the function is called and saving the address of the
30 calling instruction so that, when the function is finished, program execution can begin from where it left off. The header can also be used to perform many other useful functions at the option of the designer including initializing registers and recovering the values of saved state registers. There can also be processing to ensure consistency of processor registers.

The header 50a is often constant and can be added at the beginning of the initialization of the
35 individual machine code functions in the population. Mutation, crossover and any other operator that alters the

entity must be prevented from changing this header field when they are applied to an individual function or the field must be repaired after the application of the operator.

The footer 50c is similar to the header 50a, but does the operations in the opposite order and "cleans up" after the function call by, among other things, restoring the registers to their state before the function call, recovering the output of the function, saving state registers, and many other useful functions at the option of the designer. The footer field must also be protected from change by the genetic or other alteration operators or the field must be repaired after the application of the operator.

The return instruction 50d forces the system to leave the function and return program control to the calling procedure via the program counter 42. If variable length programs are desired, the return instruction can be allowed to move within a range defined by a minimum and maximum program size. The return instruction must also be protected from change by the genetic or other alteration operators or the field must be repaired after the application of the operator.

The function body 50b includes at least one directly executable machine code instruction, for example instructions 52a, 52b and 52c. In the SPARC architecture, instructions have a fixed length of 32 bits. Instructions can be of two types: a first type which is designated as 54 and consists of a single 32 bit operator; and a second type which is designated as 56 and includes a 19 bit operator and a 13 bit operand. The operand represents data in the form of a numerical variable.

During the initialization step, the instruction bodies 50b of the functions 50 are filled with valid machine code instructions in a suitable manner. In the preferred implementation, the instructions are initialized from a Default Instruction Template Set, which contains partially "blank" instructions that contain information on what types of instruction is represented by a particular template. For example, is it an instruction that adds the values in two registers together or an instruction that adds a constant to a register value? What is left blank in the template is the register(s) to which the instruction applies and the constant values to be included in the instruction. The register references and the constant values may be added by methods such as those described in Figure 18d. In Figure 18d, instruction templates are randomly chosen and then the blanks are filled out so that the integer array constituting the individual contains only valid machine code instructions, taking care to limit the register references therein members of the Register Set. See Figures 17c, 18b and 18d. The instructions can be selected randomly, or using any suitable selection criterion.

In the most basic form of the invention, to limit the search space and to avoid complex control mechanisms and thus achieve maximum performance, the set of machine code instructions is limited in a number of ways. Only two registers and only those machine code instructions of two addressing mode types are used. All of them operate internally in the processor. However, more lengthy and complex implementations of the Register Set are described above and more complex implementations of the instruction set are within the spirit and scope of the invention.

The first addressing mode takes one argument from memory immediately afterwards and performs an operation on this argument and a processor register and leaves the result in a register. The other takes the operand from the second register, performs the operation and leaves the result in the first register.

With these kinds of instructions it is possible to reach the maximum processor throughput. The instructions are also constrained to those which take an integer as argument and return an integer. This however does not limit the problem space to mathematical problems or just to integer arithmetic. No control instructions such as jump instructions are allowed which means that no loops can be formed. These limitations reduce the complexity and thus execution time of the individual programs.

The basic instructions are listed in the following TABLE.

	ADD (imm). ADD (reg2)	Add register one to an immediate operand or add register one to register two.
	SUB (imm). SUB (reg2)	Subtract register one from an immediate operand or subtract register one from register two.
	MUL (imm). MUL (reg2)	Multiply register one by an immediate operand or multiply register one by register two.
10	DIV (imm). DIV (reg2)	Divide register one by an immediate operand or divide register one by register two.
	OR (imm). OR (reg2)	Logical OR of register one and an immediate operand or logical OR of register one and register 2.
	AND (imm). AND (reg2)	Logical AND of register one and an immediate operand or logical AND of register one and register two.
	XOR (imm). XOR (reg2)	Logical EXCLUSIVE OR of register one and an immediate operand or logical EXCLUSIVE OR of register one and register two.
	SETHI	Set the high bits of a register, used when an operand bigger than 13 bits hits needs to be loaded into a register.
15	SRL	Logical shift right of register one to an immediate operand or logical shift right of register one a number of steps defined by register two.
	SLL	Logical shift left of register one to an immediate operand or logical shift left of register one and register two.
	XNOR	Logical EXCLUSIVE NOR of register one to an immediate operand or logical EXCLUSIVE NOR of register one and register two.

These functions can and are used to implement other common processor functions. Loading of operands into registers is performed by the OR instruction, clearing a register is performed by the AND instruction, and the negation instruction can be performed by the XNOR instruction etc.

An exemplary array 60 of functions F_0 to F_6 is illustrated in FIG. 7. Each function includes a header 50a including instructions which are collectively designated as H, a body 50b including instructions B, a footer 50c including one or more instructions F, and a return instruction 50d which is indicated as R. Although the figure shows a length of three instructions for the header and one for the footer, those particular numbers are for illustrative purposes only.

The functions of the array 60 are illustrated as having a fixed length. However, the invention is not so limited, and the functions can have variable lengths.

FIG. 8 illustrates an alternative embodiment of the invention in which only a single solution in the form of a function 62 is provided. The function 62 is configured as a continuous array that contains valid machine code instructions, and is typically larger than the functions of the array 60. Such a single array would be a typical way to implement a system not involving machine learning but involving repeated computations on run time data. Certain machine learning approaches would use such a single structure also.

After the array is initialized, it is recast as a function type array, and the registers 46 are initialized with training data. The training data is normally part of a "training set", each element of which consists of one or more inputs and a desired output. The training set represents the problem that is to be solved. However, the invention is not so limited, and the training data can alternatively include testing, validation, prediction, or any other data suitable for machine learning or for repetitive calculations on run-time data.

For the machine learning applications, the purpose of the learning process is to train the functions using the training data by causing the functions to evolve until one of them produces outputs in the output register(s) in response to the training inputs in the input register(s) that closely approximate the desired outputs. The inputs are passed to the functions by initializing the registers with the training data as described elsewhere. The functions of the array (or single function as illustrated in FIG. 8) are then executed with the training data as inputs. This execution causes the functions to produce outputs which are calculated in accordance with the instructions in the functions.

The output register or registers are then read, stored and compared with actual or desired outputs from the training data to calculate fitnesses of the functions. A function has a high fitness if its output closely approximates the desired output, and vice-versa. If the end criterion is met, the program ends.

If the end criterion is not met, the array is recast or used as a data array (e.g. integer), and the instruction bodies 50b are altered. During this process, selected individual instructions are altered, replaced, added or deleted such that the altered function includes at least one valid machine code instruction. In other words, regardless of the manner in which a particular function is altered, the program ensures that the altered function will not include any invalid instructions.

The function(s) can be altered using any machine learning algorithm. One preferred methodology is evolution of machine code structures and machine code contents using crossover and mutation operators as will

be described in detail below. However, the invention is not limited to any particular learning paradigm, nor is it limited to machine learning applications. The invention can be used, for example, to alter run-time data encoded in the machine code for the purpose of performing repetitive calculations on the run-time data.

As described above, it is important that the header, footer and return instruction will not be altered. This can be accomplished by protecting these elements from alteration using masking or other programming techniques. Alternatively, these elements can be allowed to be altered, and their initial states subsequently restored by repair or replacement. As another alternative, the instructions can be separated from the function, altered, and then returned to the function. There are other ways to protect the header, footer and return instructions from being altered, all of which are within the scope and spirit of this invention.

After performing the alteration, the program loops back to the step of recasting the array as a function and executing the function. The program iteratively repeats these steps until the end criterion is met.

The output of the system, when used for machine learning, is the function that has the highest fitness and thereby produces outputs that are closest to the desired outputs in response to the training inputs. The functions evolve as the result of the alterations, with one individual finally emerging as being improved to the highest fitness. The policy for selecting a "best individual(s)" is a matter of discretion by the designer. This system applies to any such policy.

As described above with reference to FIG. 3, the registers 46 include banks of input registers and output registers. These registers are used to pass data from a calling function to a called function, and are used to initialize the functions with training data as further described above.

More specifically, functions use their input registers to store variables. When a function calls another function, the Spare Architecture transfers the contents of its input registers to its output registers. The called function transfers the contents of the calling function's output registers into its own input registers, and operates on the data using its input registers. The opposite operation takes place when control is returned from the called function to the calling function.

Thus, the C program initializes a function with training data by storing the training data in its output registers and then calling the function. This method of initialization is preferred in the SPARC implementation, but the invention is not so limited. Other mechanisms for initializing functions with training data may be preferable with other computer architectures or for more complex application on the SPARC architecture.

An important feature of the present invention is to create and alter machine code functions as data, and execute the functions, from a high level language such as C. The following illustrates how this manipulation can be accomplished using the SPARC architecture.

As a simplified example, define a simple function that computes the sum of two input parameters and returns the result from this addition.

$\text{sum}(a,b) = a+b$

This function can be translated by a compiler in the SPARC architecture to the following five assembly language instructions.

```

save
add %i0,%i1,%i0
restore
ret
5  nop

```

These five assembly language instructions are stored sequentially in memory as five 32-bit numbers having the following decimal equivalents.

```

2178940928
2953183257
10 2179465216
2177359880
16777216

```

The "save" instruction corresponds to the number 2178940928, and is a special machine code instruction in the SPARC architecture that saves the contents of the calling function's registers and transfers the
15 parameters to the called function.

The "add" instruction adds the contents of input register "i0" to the contents of input register "i1", and stores the result in input register "i0". This instruction is coded as the integer 2953183257. Additions between other registers are represented by other integers.

The "restore" instruction restores the registers from the calling function and transfers the result from
20 the addition (present in register i0) for access from the calling function.

The return instruction "ret" jumps back to the memory location from where the function was called.

The "nop" instruction is a no operation instruction that does nothing but "entertain" the processor while jumping back. This instruction could be left out if the order of the "restore" and "ret" instructions was reversed as will be described below.

25 Implementation of the present invention involves calling a machine code function as exemplified above from the C or other high level language. Adding the numbers 2 and 3 by calling the function sum(2,3) can be represented in assembly language as follows.

```

...
mov 2,%o1
30 mov 3,%o0
call _sum
nop
...

```

These instructions are represented by these four integers:

2450530306

2484084739

2147483646

5 16777216

The first instruction stores the input parameter "2" in the output register %o1. The second instruction stores the input parameter "3" in the output register %o0.

The call instruction jumps to the location of the sum function and stores the address of itself in a register (output register %o7 by convention). This causes the program counter 42 to contain the address or
10 location of the first instruction of the sum function. In other words, it contains the address of a memory location which stores the integer 2178940928 (save), and execution continues from here until the return instruction is encountered. After the return from the function, the output register %o0 will contain the result from the summation which can be used in further calculations.

When a save instruction is executed, it copies the contents of all output registers to an unused set of
15 input registers. In this way the sets of input and output registers are used as a small internal stack for transfer of parameters and saving of register contents. The calling function stores the parameters of the called function in the output registers, the save instruction copies them to the input registers of the called function, and the restore instruction later copies all input registers of the called function into the output registers of the calling function. If the number of parameters is larger than the number of registers, the memory in the stack has to be used, as in
20 most other processors.

In this manner, the input data (training data in the machine learning system and run-time data for computational entities) is passed to the array of functions (solutions) by storing the data in the output registers of the calling function, and then calling the array.

The "nop" does nothing but keep the processor occupied while jumping to the function. In a more
25 effective but less clear version, the "mov 3,%o0" instruction is placed after the "call", which makes the processor execute this instruction "while" jumping to the function.

The following illustrates how to call the sum function defined above from the C-language.

```
typedef unsigned int(* function_ptr) ()
```

```
unsigned int sumarray[]={2178940928, 2953183257, 2179465216, 2177359880, 16777216};
```

```
30 a=((function_ptr) sumarray)(2,3);
```

These three lines of C-code will compute the sum of 2 and 3 and put the result in a variable "a" by calling a function defined by the integers in the "sumarray".

The first line of code defines the function pointer type in C, because this type is not predefined in the C-language.

35 The second line of code declares an integer array containing integers for the instructions in the sum function as defined above.

The last line of code converts the address of the sumarray from a pointer to an integer array to a

pointer to a function, and then calls this function with the arguments 2 and 3. The result from the function call is placed in variable "a".

This example illustrates the basic principles of manipulating binary programs. The sumarray can be altered like any other array as long as its contents are valid functions that can be called (executed) to evaluate the results of the algorithm.

The present method as implemented in the C programming language utilizes four instructions for initializing, altering, and executing machine code functions.

1. A first instruction that points to and designates machine code stored in a memory as data.
2. A second instruction that points to and designates machine code stored in the memory as at least one directly executable function.
3. A third instruction that alters machine code pointed to by the first instruction.
4. A fourth instruction that executes machine code pointed to by the second instruction.

Examples of the four instructions are presented below.

1. `unsigned int * theintegerpointer;`
 {declaration of integer pointer variable}
`theintegerpointer = (unsigned int *) malloc(Max_Individual_Size * Instruction_Size)`
 {The instruction creating a pointer to an integer array}
2. Alternative 1

`typedef unsigned int(*function_ptr) ();`
 {definition of the function_ptr type as a pointer to a function}
`function_ptr thefunctionpointer`
 {declares thefunctionpointer as a variable of type function_ptr}
`thefunctionpointer=(function_ptr) theintegerpointer;`
 {Instruction that typecasts theintegerpointer as a pointer to a function}
- Alternative 2

`typedef unsigned int(*function_ptr) ();`
 {definition of the function_ptr type as a pointer to a function}
`Predicted_Output = ((function_ptr) theintegerpointer)(Input Data Here);`
 {the code "((function_ptr) theintegerpointer)" is an expression within an instruction that performs a typecast.}
3. `theintegerpointer[2] = theintegerpointer[2] | 1234;`
 {XOR's the value in the second instruction with the integer 1234}
 or
`theintegerpointer[3] = 16777216;`
 {places an nop instruction in the third instruction slot, overwriting the previous instruction}

```

4.      var_a = thefunctionpointer(Input Data Here);
        or
        Predicted_Output= ((function_ptr) theintegerpointer)(Input Data Here);

```

In the CGPS embodiment of the invention which utilizes genetic crossover and mutation operations, the types of operations that can be performed depend on the type of instruction; more specifically if the instruction includes only an operator, or if it includes an operator and an operand. Examples of genetic crossover and mutation operations are illustrated in FIGs. 10 to 15.

FIGs. 10a and 10b illustrates a uniform crossover operation in which like numbers of adjacent complete instructions are exchanged or swapped between two functions. FIG. 10a illustrates two functions 70 and 72. Uniform crossover is performed by exchanging, for example, two instructions indicated as "4" and "2" in the function 70 for two instructions indicated as "5" and "9" in the function 72.

The results are illustrated in FIG. 10b. An altered function 70' includes all of its original instructions except for the "4" and the "2" which are replaced by the "5" and the "9" from the function 72. The function 72 includes all of its original instructions except for the "5" and the "9" which are replaced by the "4" and the "2" from the function 70.

FIGs. 11a and 11b illustrate "2-point crossover", in which blocks of different numbers of complete instructions are exchanged between two functions. In this case, two points are selected in each function, and all of the instructions between the two points in one function are exchanged for all of the instructions between the two points in the other function.

In the example of FIG. 11a, instructions indicated as "7" and "8" in a function 74 are exchanged for instructions "4", "7", "6", and "1" in a function 76 to produce functions 74' and 76' as illustrated in FIG. 11b.

FIGs. 12a and 12b illustrate how components of functions can be crossed over. In FIG. 12a, two instructions 78 and 80 have operators 78a and 80a and operands 78b and 80b respectively. In this case, the operator 78a of the function 78 which is indicated as "OP1" is exchanged for the operator 80a of the function 25 which is indicated as OP2. The result of the crossover is illustrated in FIG. 12b.

FIGs. 13a and 13b illustrate an example of how uniform crossover can be performed between all or parts of operands. In FIG. 13a, a function 82 has an operator 82a and an operand 82b, whereas a function 84 has an operator 84a and an operand 84b. The rightmost two bits of the operand 82b are exchanged for the rightmost two bits of the operand 84b to produce functions 82' and 84' with operands 82' and 84' as illustrated in FIG. 13b.

FIG. 14 illustrates how the operator of a function 86 can be mutated. In this case, the function 86 initially has an operator which is indicated as OP1, and is altered or replaced so that a mutated function 86' has an operator OP2. It is necessary that both operators OP1 and OP2 be valid machine code instructions in the set of instructions used by the system.

FIG. 15 illustrates how all or part of an operand can be mutated. A function 88 has an operator 88a and an operand 88b. In this example, the second least significant bit of the operand 88 is changed from "1" to "0" or "flipped" to produce an altered function 88' having an altered operand 88b'.

As set forth above, the present invention can be applied to any applicable problem by using any

suitable machine learning algorithm. The principles described in detail above can be applied to implement a particular application and computing environment.

In addition to implementing the present machine learning system on a general purpose computer as described above, it is possible to implement part or all of the system as a specially designed integrated circuit chip 5 90 such as an Application Specific Integrated Circuit (ASIC) which is symbolically illustrated in FIG. 9.

The chip 90 comprises a Central Processing Unit (CPU) 92 including a processor 94 and a RAM 96. The processor 94 includes normal CPU microcode plus microcode implementing storage, initialization, creation, and alteration operators. The RAM 96 is preferably a high speed cache memory which the processor 94 can access at a much higher speed than the processor 94 can access off-chip memory.

10 The RAM 96 can include a number of registers, or can have a conventional address based architecture. Preferably, the population of functions (solutions) is stored in the RAM 96 for rapid manipulation and execution by the processor 94. Other alternatives include additionally storing the high level program in the RAM 96. As yet another alternative, the chip 90 can include a ROM 98 for storing, for example, a kernel of the high level program.

15 The on chip memory, alternatively could be registers dedicated to storing individuals or high speed, on chip RAM that is not a cache. In addition, the CPU could alternatively execute machine learning operators such as crossover and mutation or any other operators that initialize, create, evaluate or alter individuals in microcode or in high speed ROM.

A preferred implementation of the invention evolves machine code structures and machine code 20 contents as a way of learning the solution to a problem. A detailed flowchart of this system is presented in FIG. 6 as an example for reference purposes.

The program utilizes a small tournament in combination with genetic crossover and mutation operations, and includes the following basic steps.

- 25 1. Randomly pick four individuals from the population.
2. Evaluate them in pairs, two at a time according to their fitness.
3. Let the two winners breed.
4. Replace the losers with the children of the two winners.
5. Repeat step 1-4 until the success predicate is true or the maximum number of tries is reached.

The flowchart is believed to be self-explanatory except for probabilistically choosing one operation 30 or another. These choices are analogous to flipping a coin, and can be implemented using, for example, a random number generator which generates random numbers between 0 and 1. If the number is 0.5 or less, a first probabilistic branch Pa is taken. If the number is higher than 0.5, a second probabilistic branch Pb is taken. Choices having more than two options can be implemented by dividing the range between 0 and 1 into a number of subranges equal to the number of choices.

35 FIGs. 16a to 16d are further provided for reference, and in combination constitute a detailed flowchart of the generic machine learning system of the present invention. The entries in the flowchart are believed to be self-explanatory, and will not be described in detail.

FIG. 16a outlines the general operation of the system. FIG. 16b sets forth the details of the block entitled "SYSTEM DEFINITION" in FIG. 16a. FIG. 16c sets forth the details of the block entitled "INITIALIZATION" in FIG. 16a, whereas FIG. 16d sets forth the details of the block entitled "LEARN FOR ONE ITERATION" in FIG. 16a.

5 Figures 17a through 17d are further provided for reference, and in combination constitute a detailed flowchart of the application of the invention to any computation problem that involves repeated access to run-time data. The entries in the flowchart use the terminology of this application and are believed to be self explanatory in the context of this application, and will not be described in detail.

10 Figures 18a through 18g are further provided for reference, and in combination constitute a detailed flowchart of the application of the invention as a system that induces the solution to problems by learning the structure and content of machine code entities. The entries in the flowchart use the terminology of this application and are believed to be self explanatory in the context of this application, and will not be described in detail. The mutation and crossover operators referred to in said Figure 18g are the operators described elsewhere in this application.

15 FIG. 16a outlines the general operation of the system when it is applied to a generic machine learning problem.

 FIG. 16b sets forth the details of the block entitled "SYSTEM DEFINITION" in FIG. 16a. The steps in this figure show, inter alia, what steps to take to analyze any machine learning problem to permit the designer to encode entities that contain both the Learned Elements and the Conversion Elements into a machine code entity
20 that can be created, initialized, stored, modified and executed by the means described in this application. Numeric or other values may be encoded in the machine code as constants.

 FIG. 16c sets forth the details of the block entitled "INITIALIZATION" in FIG. 16a. This Figure sets forth, inter alia, a set of steps that will result in the creation of one or more learning entity or entities that will be the learning entity or entities in the machine learning system that is being implemented using the present
25 invention. Such entity or entities will be created, stored, initialized, modified and executed by the methods set forth herein but when and how such steps take place will be set according to the requirements of the particular machine learning algorithm being implemented, as shown in FIG. 16d. It is important to note that the entity created according to the procedure outlined in Figure 16c ordinarily includes not only the Conversion Elements but also contains the first set of Learning Elements for evaluation. Should said first set not be available when the
30 initialization occurs, then the entity should be initialized with dummy values in the places in the instructions where the real Learning Elements will reside. Then, when the first set of real Learning Elements are known, it can be placed into those places using the procedure under "Start Modification" in FIG. 16d.

 FIG. 16d sets forth the details of the block entitled "LEARN FOR ONE ITERATION" in FIG. 16a. This figure shows, inter alia, how to modify and how to evaluate an entity when the particular machine learning
35 algorithm being implemented calls for either of those steps. The particular implementation of an application of the invention to a particular machine learning problem will vary substantially from problem to problem and among various machine learning systems. So this Figure is general in terms of when to evaluate the entity (referred to

in FIGs. 16a-16d as a "solution") and when to modify the entity. Because of the breadth of various approaches to machine learning, this Figure indicates that steps other than the ones shown specifically are appropriate and that systems including such other steps are within the spirit and scope of the present invention.

FIG. 17a outlines the general operation of the invention when the invention is to be used to perform
5 repeated computations on run-time data. This application of the invention could be handled in other and alternative manners both in the general matters set forth in FIG. 17a and in 17b through 17d and those other and alternative manners are within the scope and spirit of this invention.

FIG. 17b sets forth the details of the block entitled "SYSTEM DEFINITION" in FIG. 17a. This
Figure shows, inter alia, how to define a repeated computation on run-time data problem so that it may be
10 implemented using the invention. Numeric or other values may be encoded in the machine code as constants.

FIG. 17c sets forth the details of the block entitled "INITIALIZATION" in FIG. 17a. This Figure
sets forth, inter alia, a set of steps that will result in the creation of a computational entity that will perform
repeated computations on run-time data based on the system definition performed in FIG. 17b. It is important
to note that the computation entity created ordinarily includes not only the related code but also contains the first
15 set of run-time data on which the system will perform repeated calculations. Should said first set not be available
when the initialization occurs, then the computational entity should be initialized with dummy values in the places
in the instructions where the real run-time data will reside. Then, when the real run-time data is known, it can
be placed into those places using the procedure under "Start Modification" in FIG. 17d.

FIG. 17d sets forth the details of the block entitled "EVALUATE OR MODIFY FUNCTION" in FIG.
20 17a. The particular implementation of an application of the invention to repeated computations on run-time data
will vary substantially from problem to problem. So this Figure is general in terms of when to execute the
computational entity (referred to in FIGs. 17a-17d as a "solution") and when to modify the computational entity.
The computational entity should be modified each time that the designer wishes to perform a particular
computation (related code) on run-time data of the type that is encoded in the computational entity but that has
25 not yet been placed in the computational entity. The computational entity should be executed every time the
designer wants to perform a calculation on run-time data that has been encoded into the machine code of the
computational entity.

FIG. 18a outlines the general operation of the invention when it is applied to learning machine code
structures and machine code contents that will operate a register machine for the purpose of learning a solution
30 to a problem. This application of the invention could be handled in other and alternative manners both in the
general matters set forth in FIG. 18a and in 18b through 18h and those other and alternative manners are within
the scope and spirit of this invention.

FIG. 18b sets forth the details of the block entitled "SETUP" in FIG. 18a. This Figure describes
various steps necessary to setup an application of the invention to a CGPS run on a particular problem. Numeric
35 or other values may be encoded in the machine code as constants.

FIG. 18c sets forth the details of the block entitled "INITIALIZATION" in FIG. 18a. This Figure
describes, inter alia, a method for initializing a population of entities for the purpose of conducting a CGPS run.

FIG. 18d sets forth the details of the block entitled "CREATE INSTRUCTION" in FIG. 18c. This Figure describes, inter alia, one method of creating a single machine code instruction to be included in the body of an entity.

FIG. 18e sets forth the details of the block entitled "MAIN CGPS LOOP" in FIG. 18a. This Figure sets forth, inter alia, one approach to implementing the basic CGPS learning algorithm.

FIG. 18f sets forth the details of the block entitled "CALCULATE INDIV[N] FITNESS*" in FIG. 18c. The fitness function described therein is simple. More complex fitness functions may easily be implemented and are within the spirit and scope of this invention.

FIG. 18g sets forth the details of the block entitled "PERFORM GENETIC OPERATIONS . . ." in FIG. 18c. The mutation and crossover operators referred to in Figure 18g may be all or any of the operators described elsewhere in this application.

FIG. 18h sets forth the details of the block entitled "DECOMPILE CHOSEN SOLUTION" in FIG. 18a. This Figure describes a method of converting an individual into a standard C language function that may then be linked and compiled into any C application. Other and more complex methods may be easily implemented and are within the spirit and scope of this invention.

TURING COMPLETE MACHINE LEARNING SYSTEM

A basic machine learning system and method according to the present invention is described above. The system can be made Turing complete though the addition of branching instructions as will be described in detail below. This capability increases the flexibility and power of the system by enabling subroutines, leaf functions, external function calls, recursion, loops, and other types of conditional operations.

Introduction

The modern computer has grown in instruction capabilities and machine word size, which means that there are more tasks that can be carried out elegantly by a straightforward machine language subroutines. For example, the precision of the arithmetic is enough for direct usage, and the memory organization is often flat which makes pointer operations less complex, etc.

Conventional application programs often fail to use the computer to its full potential because the algorithms are not implemented in the most efficient way. This fact is even more relevant in the use of algorithms for machine learning, artificial intelligence, artificial life, etc., repeated computations on run-time data which rely on complex manipulations and are often of the meta type where programs manipulate program structures.

The present machine learning system is a unique approach to implementation of algorithms of this kind. The approach is unknown, in spite of the fact that it is possible to use on the oldest and smallest computer. It can present a number of advantages like a speed up of a 1,000 times, low memory consumption, compact representation, etc.

The execution speed enhancement is important both for real-life applications, and for simulations and experiments in science. A large efficiency improvement of this magnitude can make real life applications feasible, and it could also mean the difference of whether an experiment or simulation will take three days or one year.

This can make a given algorithm useful in a whole new domain of problems.

If an algorithm uses a representation in a program that later can be evaluated by some kind of interpretation as a program, even in a limited sense, this approach can be used. It could be the rules in Cellular Automata, the creatures in artificial life, decision trees, rules, simulations of adaptive behavior, or as demonstrated below evolving Turing complete algorithms with an evolutionary algorithm. As described above, the present approach is referred to as a compiling approach because there are no interpreting parts and the individuals are, in effect, directly compiled.

The present method of binary machine code manipulation should not be confused with translating a machine learning algorithm into assembly code, the way a compiler operates. Instead, the present invention is capable of meta manipulating machine code in an efficient way.

The Present Compiling Approach

It is surprising that there are no existing programming languages with specialized features and tools for this programming paradigm.

The present invention provides advantages of the prior art including the following.

- Higher execution speed
- Compact kernel
- Compact representation
- Uncomplicated memory management
- Low complexity algorithms

The efficiency improvement in binary meta manipulations comes from deletion of the interpreting steps. Instead of using a structure defined by the programmer to be run through an interpreter, a match between the data structures of the application and the binary machine code structures is sought. The algorithm then manipulates the binaries and instead of interpreting, the system executes the binary directly. The efficiency benefit comes mostly from deletion of the interpreting steps, but also from the simplicity of manipulating the linear integer array that constitutes the binary machine code.

The lack of an interpreting or compiling step makes the kernel of the algorithm compact, because no definition of the interpretation is needed. The machine code binary format is often compact in itself, which is important when working with applications that use a large set of structures, for example evolutionary algorithms, cellular automata, etc.

The inherent linear structure of a binary code program forces the design of a system that uses uncomplicated memory management. The decreased use of pointer and garbage collection imposes a straightforward handling of memory that is an advantage in for example real-time applications.

Some parts of the machine learning algorithm might be simplified by the restriction of the instructions to integer arithmetic. Handling of a sequence of integers is a task that a processor handles efficiently and compactly, which could simplify the structure of the algorithm.

The approach of binary manipulating machine learning algorithms builds on the idea behind a von

Neumann machine, and it is thus applicable on most of today computers from the largest super computers to the small invisible systems in cars, cameras and washing machines.

Von Neumann Machines

A Von Neumann machine is a machine where the program of the computer resides in the same storage
5 as the data used by that program. This machine is named after the famous Hungarian/American mathematician Jon von Neumann, and almost all computers are today of the von Neumann type. The fact that the program could be considered as just another kind of data makes it possible to build programs that manipulate programs, and programs that manipulate themselves.

The memory in a machine of this type could be viewed as an indexed array of integers, and a program
10 is thus also an array of integer numbers. Different machines use different maximal sizes of integers. A 32-bit processor is currently the most commonly commercially available type. This means that the memory of this machine could be viewed as an array of integers with a maximum size of $2^{32}-1$, which is equal to 4294967295, and program in such a machine is nothing more than an array of numbers between zero and 4294967295.

A program that manipulates another program's binary instructions is just a program that manipulates
15 an array of integers. The idea of regarding program and data as something different is however deeply rooted in our way of thinking. It is so deeply rooted that most designers of higher language programmed have made it impossible for the programmer to access and manipulate binary programs. It is also surprising that no languages are designed for this kind of task. There are a no higher level languages that directly support this paradigm with the appropriate tools and structures.

20 The C language is desirable for practicing the present invention because it makes it possible to manipulate the memory where the program is stored.

The Processor

The processor is a "black box" which does the "intelligent" work in a computer. The principles of different available processors are surprisingly similar. The processor consists of several parts as illustrated in FIG.
25 3, including the control logic 44 which access the memory 12, the ALU 40, and the registers 46.

The control logic 44 uses a register or some arithmetic combination of registers to get an index number or address. The content of the memory array element with this index number is then placed in one of the registers of the processor.

A register is a place inside the processor where an integer can be stored. Normally a register can
30 store an integer with the same size as the so-called word size of the processor. A 32-bit processor have registers that can store integers between 0 and 4294967295.

The most important register is the program counter (PC) which stores the address of the next instruction to be executed by the processor. The processor looks at the contents of the memory array at the position of the program counter and interprets this integer as an instruction that might be an addition of two
35 registers or placing a value from memory into a register.

An addition of a number to the program counter itself causes transfer of control to another part of the memory, in other words a jump to another part of the program. After doing an instruction the program counter is incremented by one and another instruction is read from memory and executed.

The ALU in the processor perform arithmetic and logic instructions between registers. All processors
5 can do addition, subtraction, logical "and", logical "or", etc. More advanced processors do multiplication and division of integers, and some have floating point units with corresponding registers.

Every behavior we can see in modern computers is based on these simple principles. Computers doing graphics, animations, controlling a washing machine or watching the ignition system in a car all does these memory manipulations and register operations.

10 The principles above do not anywhere prevent a program from manipulating the contents of memory locations that later will be placed in the program counter and run by the processor as it interprets it as instructions. This is the basis of the binary manipulating machine learning code in accordance with the present invention.

Machine Code and Assembly Code

Machine language is the integers that constitute the program that the processor is executing. These
15 numbers could be expressed with, for example, different radix such as decimal, octal, hexadecimal or binary. By binary machine code we mean the actual numbers stored (in binary format) in the computer.

When programming or discussing machine language, it is often impractical to use numbers for instructions. To remember that the addition instruction for example is represented by the integer 2416058368 in the SUN SPARC architecture is not natural to the human mind.

20 If we represent it in hexadecimal radix (E219922D), it will be more compact and it is more easy to deal with than its binary equivalent, but it is still not natural to remember. For this reason, assembly language has been developed. Assembly language uses mnemonics to represent machine code instruction. For example, addition is represented by the three letters "ADD".

The grammar for assembly language is very simple, and the translation or mapping from assembly
25 language to machine code is simple and straightforward. Assembly language is not, however, machine language, and cannot be executed by the processor directly without the translation step.

RISC and CISC

The present invention has been implemented using both processors of the CISC type including the Intel 80486, and of the RISC type including the SUN-SPARC architecture. CISC stands for Complex Instruction
30 Set Computer, and RISC stands for Reduced Instruction Set Computer. As indicated by the acronym, the RISC processor has fewer and less complex instructions than the CISC processor. This means that the RISC processor can be implemented differently in hardware and that it therefore will be faster.

There are advantages of both paradigms when using them for binary manipulating machine learning. The CISC has the advantages of many more types of instructions and addressing modes, and a subset can be
35 readily found which makes a particular implementation straightforward.

The RISC processor on the other hand has the advantage that the structure or "grammar" of an instruction is very simple, for example the number of bits is constant. This sometimes makes the manipulation easier. It is easier to check for a valid instruction. A RISC is also often faster than a CISC.

Although the present invention can be practiced using CISC and RISC systems, it is somewhat easier to administrate on an RISC based architecture.

Another important factor is the operating system and hardware support of the development environment. A system like UNIX which offers separate processes and hardware implemented memory protection is definitely to recommend during development, because it will save time when bugs appear during the development phase. These systems do not have to be restarted after a serious bug like an illegal instruction. It should, however, be noted that once the system is debugged it will run as safely on any platform as any other conventional program.

Structure of Machine Code Function

A procedure and a function can be regarded as very similar concepts in machine code or even in C. A procedure is a function that does not return a value. When implementing a machine learning system with a binary manipulating technique, functions are the basic structures. For example, the individuals in compiling genetic algorithms are implemented as machine code functions. A function call has to perform three different subtasks:

- Program control transfer and saving the start address
- Saving processor registers
- Transfer parameters

The most important instruction for functions and procedures is the call instruction. It is present in all kinds of processors, and works as a jump instruction that saves the memory location or address of where it was jumping from. This will allow a return instruction to return back to this memory location when execution of the function is complete.

Most call instructions save the return address in a special memory segment called the stack, but some save it internally in a register in the processor. The SUN SPARC architecture saves the return address in a register, if there is enough room for it.

A call instruction is not sufficient to make a complete function call. The contents of the registers must be saved somewhere before the actual instructions of the function are executed.

This assures that the called function will not interfere with the processing in the calling function, and gives the called function the liberty to manipulate these registers itself. The most common place to store the contents of the calling functions registers, is the stack. Some architectures like the SPARC architecture stores it inside the processor by the special save instruction.

When the execution of the function is done, the registers of the calling function have to be restored to allow this function to continue processing in the context it was working in before the function call.

The last task a processor has to accomplish in order to perform a complete function call is to transfer

the parameters to the function. It has to transfer the input parameters when the call is performed, and transfer back the return values after the execution of the function is complete. Again, this is most commonly done by storing these values in the stack, but it can also be done by using special registers inside the processor.

These four important steps: calling the function, saving and restoring registers, transfer of parameters and return of outputs can be done in different ways. It can even be done in different ways on the same processor. Normally it is the job of the compiler to produce machine code instructions that perform these three steps in a suitable way.

In order to make it possible to link object code from different compilers and to increase portability, a number of standards, recommendations and conventions have emerged. There is for example the calling convention whereby it is up to the calling function to save the registers or whereby it is up to the called function to save the registers. In the SPARC architecture, it is the called function that saves the callers registers. In a similar manner there are conventions for how parameters should be transferred.

When working with machine learning at the binary level, it is sometimes allowable to be more free, and for example not always save all registers if it is known that the called function will not use them. This not only provides a more efficient implementation, but also presents the opportunity to use special features of the architecture that are hard to express in a high level language.

As described above, the structure of a function on the machine code level can be considered as an array of integers divided having four parts:

Header
20 Body
Footer
Return

The header of a function does one or more of the three steps mentioned above. In the SPARC architecture, the header of the function saves the registers of the calling function and sometimes also transfers parameters. Which of the three steps the header does is different from processor to processor and from compiler to compiler.

Although most function structures need some kind of header to perform a function call, some special kinds of functions do not need a header, such as the leaf procedure in the SPARC architecture.

The header is fairly constant and normally does not have to be manipulated by the machine learning part of the program. It can be defined at an early stage, for example in an initialization phase of the system.

The body of the function does the actual work that the function is supposed to carry out. When the body of the function is entered, all of its parameters are accessible to it, and the registers from the calling function are saved. The body can then use any of the arithmetic instructions or call another function to compute the desired function.

35 The footer contains the "cleanup" instructions as described elsewhere.

A return instruction must always follow the footer. This instruction finds out where the call to this function was made from, and then jumps back to this location. The address to jump back to is either stored on

the stack or in special registers in the processor.

SPARC Architecture

A system that directly manipulates machine code is not easily made portable, so it is recommended to choose one platform as a base for experimentation. The preferred system includes the SPARC architecture and 5 SUN workstations. The reason for this is that is one of the most widely used architectures in the research community, with a stable UNIX operating system. It is also a relatively fast RISC architecture.

SPARC is an open architecture with several different implementations and manufacturers. The SPARC International, Inc. is a non-profit consortium for evolution and standardization of this architecture. A well known user of SPARC is SUN Microsystems Inc., which uses the architecture in all of its modern workstations.

10 SPARC Registers

As described above with reference to FIG. 3, the most important of the different kinds of registers as applicable to the present invention are the "windowed registers". It is between the windowed registers that almost all arithmetic and logic operators take place. There are a dozen other registers that are more important to the system software than to a client application. The program counter is also an important register.

15 The windowed registers are divided into four classes:

Output registers O_0 to O_7

Input registers I_0 to I_7

Local registers L_0 to L_7

There are also eight global registers G_0 to G_7 .

20 There are eight registers in each of these classes. When a save instruction is executed, it copies the contents of the output registers into a new set or bank of corresponding input registers. Register O_0 is copied into I_0 , O_1 is copied into I_1 , etc. The input register used is a new input register owned by the called function.

The values in the old input registers are kept or saved. It will be noted that the contents are not really copied. In reality these registers are the same, and only a pointer to the current set of register is incremented.

25 The processor has an internal storage for a few banks or sets of register like this (seven in the SPARC architecture), and if this limit is exceeded, the hardware and system software will save the contents in memory. For the user, this banked register mechanism can be thought of as a small internal stack.

It is important to note that a save instruction copies the contents of the calling function's output registers into the called function's input registers, while the restore instruction copies the contents of the returning 30 function's input registers into the calling function's output registers.

When a function wants to call another function and pass some variables, it thus places the parameters in its output registers and makes the call. After the function call has been completed, the returning values can be found again in the output registers.

The local registers are local to the function working for the moment, and are used for temporary 35 storage. A fresh set of local registers is provided when a new function is called, but there is no special transfer

of values into these registers.

The global registers are always the same. They keep their meaning and content across function calls, and can thus be used to store global values. There are no alternative sets of these registers like there are with the input, output, and local registers.

5 Register Implementation

Some of the registers have a reserved usage. In some cases the reserved meaning is due to hardware, but in most cases the origin of these constraints are software conventions, as specified by the SUN Application Binary Interface (ABI).

In accordance with the present invention, almost all of the actual algorithm is written in a high level language like C. This prevents interference with the work done by the compiler. If this goal is abandoned, and the complete system were written in assembly language, even higher flexibility could be achieved. However, all portability would be lost, and it is therefore preferred to provide the manipulating program in a high level language such as C.

The global registers are preferably not used by the program, because these registers are likely to be used by the code generated by the compiler. Global storage for function structures can be provided in memory.

Global register zero has a special meaning. It is not a register where values can be stored. An attempt to read global register zero, always returns the value zero, and an attempt to write to global register zero does not change anything. This register is used when a zero constant is needed or when the result of an operation should be discarded.

Global register one is by convention assumed to be destroyed across function calls, so this register can be used by the functions that are manipulated by the machine learning algorithm.

Registers I_6 , I_7 , O_6 , and O_7 are by convention used to store stack and frame pointers as well as the return address in a function call, so these should not be used by the program.

Local registers zero and one have special uses during an interrupt but this will not affect the functions.

Thus, the registers which are available for use in the SPARC architecture are global register G_1 , input registers I_0 to I_5 , output registers O_0 to O_5 , and local registers L_0 to L_7 .

SPARC Instructions

SPARC is a RISC architecture with a word length of 32 bits. All of the instructions have this size. Basically there are three different formats of instructions, defining the meaning of the 32 bits. The processor distinguishes between the formats by looking at the last two bits, bit 30 and bit 31. The three formats are:

CALL instruction

Branches, etc.

Arithmetic and logic instructions

In the CALL instruction, bit30 is one and bit31 is zero. The rest of the bits are interpreted as a constant that is added to the program counter (PC). The return address is stored in output register I_7 .

Branches are mostly used for conditional jumps. They look at the last performed arithmetic operation, and if it fulfills a certain criteria, for example if the result is zero, then the program counter is incremented or decremented by the value of the last 22 bits in the instructions. In a branch, both bit 30 and bit 31 are zero.

The last group of instructions are the arithmetic and logic instructions between registers. These 5 groups of instructions have bit 31 as one and bit 30 as zero. These are the preferred instructions for practicing the present invention. These instructions perform, for example, multiplication, division, subtraction, addition, AND, OR, NOT, XOR and different SHIFT instructions.

The arithmetic instructions can also be used as jump, call, and return instructions, if the result of the operation is put into the program counter register. In this way it is possible to jump to the address pointed to by 10 the contents of a register. When the current value of the program counter is saved in another register (out7), these instructions can be used as call instructions.

Call is equivalent to `jmp1 register,%o7`, where `jmp1` means "jump long"; jump to the content of "register" and place the current program counter in register `O7`.

Return (`ret`) from a normal procedure is equivalent to `jmp1`.

15 Return from a leaf procedure (`ret`) is equivalent with `jmp1`.

The constant eight causes the control to jump back past the original call and past this call's delay slot.

The return instruction puts the value of `O7` with a constant added thereto into the program counter, and causes the execution of the program to return.

There are also instructions for loading and saving information from the memory.

20 Control transfer instructions (CTI) like jump, call, and return, are somewhat special in the SPARC architecture. The instruction immediately after the jump instruction is executed during the transfer of control, or it could be said to be executed before the jump.

A SPARC assembly code listing can appear misleading because of the change in execution order, because a NOP instruction is placed after the CTI. The instruction after the CTI is say to be in the delay slot of 25 the CTI. The reason for this somewhat awkward mechanism is that if the delay slot can be filled with a useful instruction, it will make the execution of the overall program more effective, because the processor can do more instructions in the same cycles. It is the hardware construction of the processor that makes this arrangement necessary.

The CTI and the delay slot are important in implementing the CGPS. Special rules require that a CTI 30 that has a CTI in its delay slot. This is called a Delayed Control Transfer Couple.

The execution of a DCTI couple is awkward. The first (and only the first) instruction of the function that the first CTI is pointing to is executed, then directly the execution jumps to the function that the second CTI is pointing at. The instruction after the second CTI is not regarded as being in a delay slot, and it is executed when control is returned. This phenomena is important for the construction of subroutines and external functions.

Leaf Procedures

The above described properties of the call and save instructions make it possible to use two particular kinds of procedures.

The first kind of procedure is a full procedure that uses the save and restore instructions, and the procedure is consequently allowed to use and manipulate input, output and local registers. The save and restore functions, however, consume processor cycles. If there is room inside the processor, the time consumption is moderate, but if storage in memory is needed it will take many processor cycles.

The solution to this problem is to use leaf procedures. They are called leaf procedures because they cannot call another procedure, and therefore leaves in the procedure structure of a program.

10 A leaf procedure does not perform a save operation, and works with the same set of registers as the calling procedure. To avoid interference with the content of the calling procedure, it only manipulates the output registers, which are assumed to be destroyed by the compiler across function calls. One of the elegant consequences of this technique is that the calling procedure does not have to know what kind of procedure it is calling. This means that linking of procedures works normally.

15 The difference between a leaf procedure and a full procedure is that it only manipulates the output registers, does not use save or restore, and has a special return instruction that looks for the return address in output register O₇ seven instead of input register I₇.

These details originate mostly from software conventions implemented in compilers. When the underlying principles are known, it is possible to have many more combinations of behavior of procedures. In 20 some of examples in *CPGS, hybrids of the normal procedures and leaf procedures are used, because it is known how the registers are used in the calling procedures.

Although the floating point features of the computer architecture are not specifically addressed herein, it is possible to use binary manipulating algorithms for floating point processors and units. This makes the present invention applicable to new domains of problems.

25 Evolutionary Algorithms

Examples of genetic algorithms were present above with reference to FIGs. 10 to 15. A genetic algorithm, for example, is an algorithm based on the principle of natural selection. A set of potential solutions, a population, is measured against a fitness criteria, and through iterations is refined with mutation and recombination (crossover) operators. In the original genetic algorithm, the individuals in the populations consist 30 of fixed length binary strings. The recombination operators used are the uniform and 2-point crossover operators.

In genetic programming the goal of the system is to evolve algorithms or programs in a given language. Most genetic programming systems use a tree structure for the representation of programs, as in the above referenced patent to Koza. The most commonly used tree representation form is the S-expressions used in the LISP-language.

35 The tree representation guarantees that the evolved programs will be syntactically correct after the genetic operators are applied. In the original model, the only genetic operator apart from selection is the subtree

exchanging crossover. This crossover operator swaps two subtrees in the individuals that undergo crossover.

A middle way between these two representation forms is messy genetic algorithms, which has a freer form of representation where for example the loci of gene is not tied to its interpretation. Instead genes are tagged with a name to enable identification.

5 Evolution of Machine Code Structures and Content

A basic machine learning system was described above with reference to FIGs. 1 to 18. Our discussion of the prior system provides many of the basic details of the present invention and is incorporated here by reference. The algorithm evolves fixed length binary strings, and uses crossover and mutation operators. The crossover is prevented from damaging certain bit fields within the 32 bits of the instructions. This protection procedure guarantees that only valid machine code instructions are generated through crossover. The mutation operator has a similar protection feature.

The system described below is a more complete machine learning and induction system, capable of evolving Turing complete algorithms and machine code functions. The system provides additional advantages including the following.

- 15 Use of several machine registers
- Dynamic allocation of memory, (no recompilation needed)
- Variable length of programs.
- Multiple input parameters to functions
- Unlimited memory through indexed memory
- 20 Automatic evolution of subfunctions
- If-then-else structures
- Jumps
- Use of loop structures including for, while, repeat
- Recursion, direct and through subfunctions
- 25 Protected functions e.g. division
- String functions, and list functions
- Linking any C-function for use in the function set

These goals are met by a program written mostly in C, which uses unrestricted crossover at instruction boundaries. Unrestricted crossover means that the crossover acting on the strings of instructions should be able to work blindly without checking what kind of instructions are moved and where.

There are many advantages of this if the goal is met. The implementation is very efficient, because the algorithm will only consist of a loop moving a sequence of integers, something a computer is very good at. The implementation will be simple and easily extendable because there will be a minimum of interaction and interference between parts of the program. Equally important is the fact that the program will be more easily ported to different platforms because the architecture specific parts can be restricted to a minor part of the program, and the crossover mechanism does not have to be effected.

It is easy to find examples of instructions and combinations of instructions where these properties do not hold. The normal call instructions constrains an offset that is added to the program counter. If this instruction is moved to another memory position, the call will point to another memory location where there might not be anything like a proper function. The SPARC architecture does not have calls with absolute addresses, which would still work after crossover. Instead, a call to an address specified by a register is used. The value in the register will be the same even if the instruction is moved by crossover.

The same problem exists with a normal branch instruction. It is also defined by a constant added to the program counter, and it cannot either be moved without changing this constant. But a certain class of instructions are possible to move unrestrictedly, without any unwanted effects. Machine language is quite flexible, and this restriction of instruction use is possible without abandoning the goals of the system.

Crossover

The crossover operators (uniform and 2-point) were described in detail above. See Figure 22j.

Mutation

The mutation operator picks an instruction at random, and checks whether it has a constant part or if it is only an operation between registers. If it has a constant part, a bit in this constant is mutated and also potentially the source and destination registers of the operation. If the instruction does not have a constraint part, the instruction's type, source and destination registers are mutated. See Figure 22i.

Calls and jumps are not mutated other than that they may be swapped for other instructions.

Casting and Execution of an Array

In a computer, all data is represented as integer numbers. To make programming a more easy task there are more data types and structures in a high level program, for example: strings, pointers, characters bitmaps, etc.

But all of these structures are, as stated above, translated into integers. Sometimes the programmer wants to translate an object of a certain type into an object of another type. The term for this is casting, one type is cast into another type. This is often only an operation in the high level language, at the machine code level they might be represented by the same integer.

Speed Increase

The efficiency of the present invention can be 1,000 times faster than a system coded in an interpreting language, and there are reasons to believe that similar performance enhancements are possible for other kinds of meta manipulating programs. If, however, these speed enhancements still would not be enough there are a few ways to further improve performance.

The main system described here has a number of advanced features like subroutines, recursion, etc. If extremely fast performance is needed and if a system can do without these features and if the goal is to only

evolve arithmetic and logic expressions of a number of "in" variables, then the efficiency can be increased even further.

More specifically, if it is acceptable to decompile the system between runs, for example if the problem parameters are given beforehand, then execution speed can be increased further. Finally, there is the possibility of coding and optimizing the entire system in assembly language which has the fastest execution.

System Parameters

The following is a list and a brief description of the parameters that are used to control the present machine learning system.

- Population size
 - 10 Number of fitness cases to evaluate. Training will halt when the system has tried these many fitness cases.
 - Mutation probability
 - Crossover probability
 - Parsimony pressure start value
 - 15 Number of incremental steps for parsimony pressure
 - Size of each incremental parsimony pressure step
 - Maximum program size
 - Initial mean program size
 - Random number generator seed
 - 20 Success threshold for fitness
 - Maximum number of iterations, in recursion and loops
 - Maximum size of initial terminal set numbers
 - Number of bits to be mutated in terminal numbers
 - Number of input registers to use
 - 25 Number of output registers to use
 - Number of external functions
 - Maximum number of automatic subfunctions
 - Flags determining which instructions to use: ADD, SUB, MUL, SLL, SRL, XOR, AND, OR
- These parameters are read from an initialization file before each training session. All effects of these
- 30 parameters are dynamically allocated, and there is no recompilation of the system necessary. This implies that a flexible system can be implemented on a small computer without an onboard compiler. These parameters would, of course, be different for other machine learning systems or for repeat calculation systems on run-time data.

Machine Code Instruction Implementation

- The following SPARC machine code instructions are used in the CGPS implementation.
- 35 ADD, Addition

- SUB, Subtraction
- MUL, Integer multiplication
- SLL, Shift left
- SRL, Shift right
- 5 XOR, Exclusive or
- AND, Logical And
- OR, Logical Or
- Call $L_{0..7}$

The arithmetic and logic instructions, all instructions except the last call instructions, come in four
 10 different classes. These arithmetic instructions can have the property of affecting a following if-then branch or not affect it. They can also have the property of being an operation between three registers, or an operation between two registers and a constant. The combinations of these two classes makes four different variants of the arithmetic instructions.

An arithmetic instruction could for example add output register O_1 with output register O_2 and store
 15 the result in output register O_3 , or it could add a constant to output register O_2 and store the result in output register O_3 .

The 32-bit instruction format has room for a constant of ± 4196 . In this manner, a single instruction is substantially equivalent to many elements in an ordinary machine learning system: one element for the operator, two elements for the two operands, and one element for the destination of the result. This approach is thus quite
 20 memory effective using only four bytes of memory to store four nodes.

There are eight different call instructions, one for each of the local registers. The call instruction jumps to an address given by the content of the local registers in the function. The local registers are thus reserved for this usage in our implementation, and they have to be initialized in the header of every function to make sure that they point to a function in use.

25 These instructions are the basis for external functions and subroutines. How many of these instructions are used in the initialization is determined by how many subroutines and external functions that are used.

The maximum number of external functions and subroutines are in this implementation limited to eight. If more external functions are needed, then another initialization technique can be used where the call
 30 instructions jumps to the address given by a constant added to a local register. The instruction format allows for storage of such constants within the instruction.

The division instruction can be used directly as a basic instruction. Protection from division by zero can be provided by catching the interrupt generated by hardware. It is, however, more effective to make protected division as an external function, with a few instructions checking for division by zero.

Initialization

Initialization is the first task the system performs when presented with a giving training situation. See FIGs. 22b, 22c. Initialization can be divided into four steps.

- Allocation of memory for the population, etc.
- 5 Initialization of header
- Initialization of footer
- Initialization of the function body

Initialization of Memory FIGs. 22b, 22c.

The memory used for the individuals in the arrays is a linear array of integers. The array is divided
 10 into blocks determined by the system parameter maximum length. A fixed maximum length is thus reserved for every individual. If there are subroutines, then this memory is allocated for every subroutine according to the maximal number of subroutines. The program and its subroutines then can vary in length within these boundaries.

The advantages with this paradigm include very simple memory management, without garbage collection. The approach with linear memory is efficient and natural for the use of binary code. There is also
 15 a constant -- easily calculated -- memory usage once the system is initialized. This can be a requirement in real-time systems, and in systems with limited memory resources.

An example population of one million individuals with no subroutines and 40 nodes maximum per individual in this way occupies 40MB of RAM memory, which is manageable on an ordinary workstation.

Initialization of Header FIGs. 22b, 22c.

20 The initialization of the header consists of two parts; one that is fixed, and one that depends on the number of subroutines and external functions.

The fixed part of the header is a NOP instruction and a save instruction.

The NOP instruction is needed for the rare case that this function is called from another function as a subroutine by a control transfer couple as described above. A control transfer couple can arise after an
 25 unrestricted crossover where two jumps or call instructions are placed after each other.

Only the first instruction of the first call is executed. If this first instruction is a save instruction, which is normal, then this save would be executed alone and the control will go to the address of the second call which probably will also be a save instruction.

The first of these two save instructions will be unbalanced, not corresponding to a restore, and the
 30 registers will be corrupt. If instead a NOP is placed in the header of every instruction, the NOP instruction can be executed safely without affecting the state of the machine.

This point is probably the most special case in the implementation, but it enables the use of unrestricted crossover. When an external function is to be used in the system, it should also have a NOP instruction or another harmless instruction as its first instruction.

35 The second part of the header initializes the local registers. The local registers are used to store jump

addresses of subroutines and external functions. This part of the header contains load instructions that load the appropriate addresses into the local registers. The current addresses are put here in the header during initialization, and then when an individual function is executed, it first executes the header and thus puts the right values in the local registers, which guarantees that the later call and jumps will be performed to the desired addresses. This
 5 part of the header varies in sizes by the number of subroutines and external functions.

Initialization of Footer FIGs. 22b, 22c.

The initialization of the footer is simple. The footer consists of two fixed instructions; one NOP instruction followed by a restore instruction. The NOP instruction is used in the last instruction in the body as a control transfer instruction like a jump or a call. In this case, the NOP instruction goes into the delay slot of
 10 the jump instruction and is executed during this procedure call. If the NOP instruction was not present, the restore instruction or the return instruction that follows the footer would go into the delay slot which would corrupt the registers. We have previously discussed headers and that discussion applies here also.

Initialization of Return Instruction

A return instruction must follow the footer.

15 Initialization of Function Body FIGs. 22b, 22c, 22d.

The function body is initialized by for each memory cell by randomly selecting an instruction from the set of instructions that the user has put into the function set, including call instructions using local registers. If the picked instruction is an arithmetic instruction, input and output registers are chosen for operands and destination according to the parameters supplied by the user.

20 We have previously discussed headers, bodies, footers and return instructions, and that discussion applies here also.

With a certain probability, an instruction is given either a constant and register as operands, or two registers as operands. If one of the operands is a constant, this constant is randomly generated to a maximum size defined by its parameter and put into the instruction. The instruction has room for constants within the range of
 25 ± 4196 , within the 32 bits of the instruction.

Subroutines

Subroutines are modularisations within an individual that spontaneously change during evolution. A subroutine has the same structure as a function call, and consists of parameters transferred in the out variables, and a call function.

30 An individual in this system is a linear array of numbers. This array is divided into a number of pieces of uniform size. The number of pieces corresponds to the maximum number of the subroutine parameter. Every such memory section is a subroutine. A subroutine is organized the same as a main function, with a header, a footer, and a function body.

In the header of a subroutine, the local registers are initialized to contain the addresses of the other subroutines that can be called from this subroutine. When the recursion option of the system is switched off, the local registers are only loaded with the addresses of subroutines higher up in the hierarchy. So if the maximum number of subroutines is set to four, the local registers L_0 to L_3 in the main function are initialized with the addresses of subroutines 0 to 3, while the local registers L_0 and L_1 in the first subroutine are initialized with the addresses of subroutines 3 and 4.

Remaining unused local registers are initialized with the value of a dummy function that executes a return when called. With the scheme it is possible to allow unrestricted crossover between individuals and between subfunctions, because the local registers will always be initialized in the header of each subfunction to a correct address. The "call local register" instructions can thus be freely copied in the population.

Recursion

Recursion can be implemented by initializing the local registers not only to the values of subroutines higher up in the hierarchy, but also to the current subroutine itself. Recursion can also be implemented in a small external function, a leaf function. The difference between the two approaches is small, but the main advantage of the later method is that the header can be kept the same regardless of if recursion is used or not, which makes the implementation less complex.

Regardless of which approach is used for recursion, there will always be the problem of infinite chains of instruction calls. The halting problem makes it impossible to know in advance which recursive functions will stop and which will not.

The solution to this is to have a global variable in a memory location that is incremented every time a function is called. If a certain limit is reached, the execution of the individual is abandoned. The code for this checking is placed in the header of each function if the first recursion method is used. If the second method is used, this code is the first part of the recursion leaf procedure.

Leaf Procedures as Program Primitives

Loops are implemented in a way similar to recursion. A leaf procedure is used which performs a test of a variable. Depending on the outcome of the test, a loop branch is either performed or not.

The test can be whether the last performed arithmetic instruction produced zero as a result. This is accomplished by checking the so called zero flag in the processor. Out loop structures can be used simultaneously by checking other integer conditions from the last instruction.

This branch is made with a return instruction, which is a "longjump" instruction jumping to the address of a register with a constant added thereto. This constant can be positive or negative. The normal return instruction jumps back to the address given by the content of register O_7 or I_7 , incremented by eight. These eight bytes cause the return to skip the original call instruction and its delay slot.

When this function is used in a loop, the constant is made negative so that return jumps back to an address before the call to the "loop leaf procedure". The fact that the different call addresses so far are in the

execution of the individual, all are stored and accessible to the leaf procedure which makes it possible to construct a large number of control structures. With this method, it is possible to define efficient external leaf procedures that implement ordinary jumps, all kinds of loops, indexed memory, if-then-else structures, and protected functions.

A limit of eight external functions is implied by the number of local registers. This is enough for many applications, but if other domains demand more external functions and subroutines, it is possible to define their addresses as constants added to local register. This technique provides the possibility of using an almost unlimited number of primitives, external function and subroutines.

External Functions

It is possible to incorporate any C-function into the function set. Any C-module can be compiled and linked into the system. There are a number of steps that must be taken in order to make this linking successful.

First the module could be compiled to assembler code by the "-S" flag of the "cc" compiler. A NOP operation is then added before the other instructions in the function. The name of the function is added in a call array in the main C kernel code, and potentially a string name is added to an array for disassembling. After recompilation the system is ready to use the new external function. This approach can be extended and made more automatic with dynamic linking, etc.

C-Language Output

A disassembler can be provided which translates the generated binary machine code into C-language modules. This disassembler feature is provided with the goal that the output from the system should be able to be used directly as a C-module, and it should be possible to compile and link it to another c-program. With this feature the system can be used as a utility to a conventional development environment. The disassembler could also be requested to produce assembler code.

Portability Methods

The main disadvantage with the machine code manipulating technique is that it is machine specific. There are, however, a number of methods to make the system more portable.

The first is to allow for unrestricted crossover. This will make the machine independent part of the code quite small, only defining a number corresponding to instructions, etc. In this way it is possible to write a system that runs on multiple platforms without changes in the code.

An more radical approach is to let the system look for itself in its binary code how it is translated to machine language. Many template functions exist in the C-code with translations to machine code that define how the processor is working. This method is only possible for compiling systems that do simple functions with arithmetic, without jumps, subroutines, etc.

Another approach is to lean more on the hardware interrupts for portability. For example, every time an illegal instruction is encountered by the processor, a hardware interrupt is generated. Using the interrupt features in Unix it is possible to be less restrictive when manipulating binaries knowing that the system will catch

some of the erroneous structures.

This increases portability, because faults can sometimes be ignored that arise due to incompatibilities of different architectures. In an extreme example, a machine learning system can be implemented by using mutation and crossover at the bit level where all of the many illegal situation are caught by the processor memory management hardware or file system protection mechanisms.

The ideal portability situation would be to have a special language, present on different platforms, for this kind of run-time binary manipulation.

Using Tree Representation

In the preferred implementation of the present invention, a binary string approach is used for representation of individuals. Binary strings are the common representation form of genetic algorithms while tree representations are used in conventional genetic programming.

There may be applications where tree representation and crossover at a tree level is interesting to study and consequently it would be interesting to have the individual in a tree form. This can be accomplished in several ways.

One way is to use two different NOP instruction as parenthesis in the code. The crossover operator then scans through the code, matches these parenthesis, and performs tree crossover accordingly. This could potentially make the crossover operator less efficient for large individual sizes.

Another way is to break up the individual into segments in the memory, where part of the tree is presented as jump statements to subtrees. This would still make it possible to execute the individual directly. The crossover operator here must also scan through the individual, but it only has to move smaller blocks of code that could be allocated dynamically.

The third way is to let every individual have an extra array associated with it, that carries information about where subtrees start and stop. This array will not be involved in the execution of the individual, and it will only be used by the crossover operator to locate subtrees.

25 Applications

The present invention can be applied to any problem in which a computer algorithm manipulates a structure that later should be interpreted as instructions. Examples of these kind of such applications include the following.

- Genetic algorithms and genetic programming
- Cellular automata
- Artificial life systems like Tierra
- Rule induction systems
- Decision trees
- LISP or PROLOG interpreters
- 35 The invention is especially suited for applications within areas that require:

- High execution speed
- Real time learning
- Large memory structures
- Low end architectures, e.g. consumer electronics
- 5 Well defined memory behavior

Example

FIGs. 19 to 21 illustrate a function structure and associated registers in accordance with the Turing complete machine learning system of the present invention.

FIG. 19 illustrates an array of functions F_0 to F_6 , each of which consists of a main function MAIN 10 and two subroutines SUB1 and SUB2. Each function has a maximum, but variable length. Portions of the functions occupied by the header are indicated as H, the instruction body as B, the footer as F, and the return instruction as R.

It will be noted that the main functions and the subroutines SUB1 and SUB2 also have variable, but maximum lengths. Any memory not occupied by an operative portion of a function is unused.

15 FIG. 20 illustrates a function FUNCTION which consists of a main function MAIN, and two subroutines SUB1 and SUB2. Further illustrated are the input registers, output registers and local registers of BANK0 which is used by the main function, a leaf function, and a dummy function. The latter functions are stored in the memory 12.

The starting addresses of the functions MAIN, SUB1 and SUB2 are designated as ad0, ad1, and ad2, 20 whereas the starting addresses of the functions LEAF and DUMMY are designated as ad3 and ad4 respectively. The instructions which can be placed in the functions MAIN, SUB1 AND SUB2 and which are subject to alteration for the purpose of machine learning are limited to those which branch to the addresses ad0 to ad4.

The function LEAF performs the operation of protected division, where a variable "c" stored in the input register I_2 is to be divided by a variable "b" stored in the input register I_1 . More specifically, an instruction 25 "TEST I_0 " tests whether or not the variable "b" is zero. If so, the result of the division would be an infinitely large number, which constitutes an error for the system.

If the test instruction detects that the value of I_1 is zero, the next instruction is skipped and the next instruction, which is, a RETURN instruction, is executed, returning control the calling function.

If the test instruction detects that the value of I_1 is not zero, the division operation c/b is performed, 30 the result stored in the input register I_3 (as a variable d), and control returned to the calling function.

The function DUMMY consists of a return instruction that merely returns control to the calling function.

The headers of the functions MAIN, SUB1 AND SUB2 each include a SAVE instruction, and three instructions that initialize the local registers with the addresses of functions that can be called by the particular 35 function. The SAVE instruction causes the contents of the output registers of a calling function to be copied into the input registers of the called function as described above.

In the function MAIN, these initialization instructions cause the addresses ad1, ad2 and ad3 to be stored in the local registers L_0 , L_1 , and L_2 respectively.

In the normal hierarchical arrangement of the invention, functions are only allowed to call functions which are lower on the hierarchy, and leaf functions. Thus, the function SUB1 is allowed to call the function SUB2, but not the function MAIN. This is accomplished by storing ad2 for the function SUB2 in L_0 , ad3 for the function LEAF in L_1 , and ad4 for the function DUMMY in L_2 .

The function SUB2 is only allowed to call leaf functions. Therefore, the address ad3 for the function LEAF is stored, in L_0 , and the address ad4 for the function DUMMY is stored in L_1 and L_2 .

FIG. 20 further illustrates how an arithmetic instruction is performed by executing a function, and how a variable is passed from a calling function to a called function.

More specifically, execution of the instruction $I_0 = I_1 * 2$ by the function MAIN causes the variable "d" in the register I_1 to be multiplied by 2, and the result stored in the register I_0 as indicated by an arrow. The contents of the register I_0 are indicated as $a = b * 2$.

To pass the value of "a" (in I_0) to a called function, an instruction $O_0 = I_0$. This causes the contents of the input register I_0 to be copied into the output register O_0 . The next instruction, CALL L_0 , causes control to be transferred to the function having the starting address ad1, which is the function SUB1. The call instruction causes the contents of the output registers of the function MAIN to be copied into the input registers of the function SUB1, and thereby pass the value of the variable "a" which was stored in O_0 of the function MAIN to the input register I_0 of the function SUB1.

Upon return from the call to SUB1, the contents of the input registers of the function SUB1 are copied to the output registers of the function MAIN. The next instruction in the function MAIN is $I_0 = O_0$. This causes the contents of O_0 to be copied to I_0 , and thereby pass the variable that was generated by the function SUB1 and stored in I_0 thereof to the function MAIN.

FIG. 21 is similar to FIG. 20, but illustrates an arrangement including one subroutine SUB1 and one leaf function, with the subroutine SUB1 being allowed to perform recursion. In this example, the function SUB1 can call the function MAIN, itself, and the function LEAF.

This is accomplished by storing ado (the address of the function MAIN) in L_0 , storing ad1 (the address of itself) in L_1 , and ad2 (the address of the function LEAF) in L_2 .

A detailed flowchart of the Turing complete machine learning system is illustrated as a flowchart in FIGs. 22a to 22k.

FIG. 22a is main diagram of the system.

FIG. 22b illustrates the details of a block SETUP in FIG. 22a.

FIG. 22c illustrates the details of a block INITIALIZATION in FIG. 22a.

FIG. 22d illustrates the details of a block INITIALIZE LOCAL REGISTERS IN HEADER in FIG.

22c.

FIG. 22e illustrates the details of a block CREATE INSTRUCTION in FIGs. 22c and 22i;

FIG. 22f illustrates the details of a block MAIN CGPS LOOP in FIG. 22a.

FIG. 22g illustrates the details of a block CALCULATE INDIV[N] FITNESS in FIG. 22f. The steps indicated by a left bracket are performed by a C-instruction "Predicted_Output[i]=((function_ptr) Indiv[n]0(input[1]...Input[k]))".

FIG. 22h illustrates the details of a block PERFORM GENETIC OPERATIONS in FIG. 22f.

5 FIG. 22i illustrates the details of a blocks MUTATE INDIV[1] and MUTATE INDIV[3] in FIG. 22h.

FIG. 22j illustrates the details of a block CROSSOVER INDIVS [1] and [3] in FIG. 22h.

FIG. 22k illustrates the details of a block DECOMPILE CHOSEN SOLUTION in FIG. 22a.

LEARNING AND CONTROL METHOD AND SYSTEM EXEMPLIFIED
BY A METHOD AND SYSTEM THAT LEARNS HOW TO CONTROL AN
 10 AUTONOMOUS AGENT SUCH AS A ROBOT

The present invention comprises a system and/or method to perform any learning task. In its exemplary embodiments presented herein, the present invention comprises a general method for performing real time or online learning tasks such as control of an autonomous agent or a robot. In the exemplary embodiments presented below, the autonomous agent is a propelled object, more specifically a robot. However, it will be
 15 understood that the invention is not so limited, and the present invention can be applied to any applicable learning task, any applicable real time or online learning task, or to control any applicable autonomous object or process.

Additional disclosure relating to this aspect of the invention is presented in a paper entitled "REAL TIME EVOLUTION OF BEHAVIOR AND A WORLD MODEL FOR A MINIATURE ROBOT USING GENETIC PROGRAMMING", by Peter Nordin et al. Department of Computer Science, University of Dortmund,
 20 pp. 1-32, November 1995, which is incorporated by reference herein in its entirety.

Two preferred embodiments of this aspect of the present invention are described herein. The first is the **non-memory** method or system. The second is the **memory** method or system.

The two preferred implementations of this aspect of the present invention both utilize a fitting algorithm that is able to derive a function or program that takes one set of data as input and predicts the value
 25 of another set of data. The particular fitting algorithm used in the preferred embodiments of the present invention is a symbolic regression algorithm, preferably the Compiling Genetic Programming System (CGPS) as described in detail above running a symbolic regression algorithm. Symbolic regression is discussed in more detail below.

The non-memory embodiment of the present invention includes the following elements:

1. A population of solution candidates (programs), where population size can vary between 30
 30 and 50,000 individuals. The population is normally initiated to a random content. The solution candidates are referred to elsewhere in this application as "individuals," "entities," or "solutions."
2. A fitness measure defining the desired task of the evolved programs.
3. A recombination or crossover operator allowing for exchange of solution segments.
4. A mutation operator for changing code more locally than the crossover operator.
- 35 5. A set of basic constants, instructions, procedures or functions used as the atomic parts of the evolved programs.

Genetic Programming and Machine Code

The Genetic Programming (GP) system used in the preferred embodiment of the present invention to perform symbolic regression is a variant of GP that uses a linear genome, and stores the individuals of the population as binary machine code in memory. This results in a speed-up of several orders of magnitude. The method is also memory efficient requiring only 32KB for the GP kernel. The individuals can be stored in an economic way, and memory consumption is stable during evolution without any need for garbage collection etc. All of these properties make the system ideally suited for real-time control in low-end processor architectures such as one-chip embedded control.

The present CGPS system uses variable length strings of 32 bit instructions for a register machine. Each node in the genome is an instruction for a register machine. The register machine performs arithmetic operations on a small set of registers. Each instruction might also include a small integer constant of maximal 13 bits. The 32 bits in the instruction thus represent simple arithmetic operations such as "a=b+c" or "c=b*5". The actual format of the 32 bits corresponds to the machine code format of a SUN-4, which enables the genetic operators to manipulate binary code directly.

The set-up is motivated by fast execution, low memory requirements and a linear genome which makes reasoning about information content less complex. This compact system is a prerequisite for a micro-controller version.

The machine code manipulating GP system uses two-point string crossover. A node is the atomic crossover unit in the GP structure. Crossover can occur on either or both sides of a node but not within a node. Because our particular implementation of GP works with 32 bit machine code instructions, a node is a 32 bit instruction. Mutation flips bits inside the 32-bit node. The mutation operator ensures that only instructions in the function set with valid ranges of registers and constants are the result of a mutation. All genetic operators ensure syntactic closure during evolution.

The instructions are all low-level machine code instructions. An exemplary function set consists of the arithmetic operations ADD, SUB and MUL, the shift operations SLL and SLR, and finally the logic operations AND, OR and XOR. All these instructions operate on 32-bit registers.

The following illustrates how an individual program may look if printed as a 'C' program.

```
a=s3 + 4;
d=s2 >> s1;
30 b=s1 - d;
d=s2 + 2;
c=d >> 1;
b=d - d;
a=c - 3;
35 c=s4 << a;
d=a * 4;
a=a ~ 5;
```

```

c=d + s4;
c=b & c;
d=d + d;
c=d | 8;
5 d=d * 10;
b=c >> 6;
d=b & 0;
c=c >> b;
motor2=a | c;
10 c=d | 7;
motor1=c * 9;
c=e & e;
(">>",">>"]="Shift "|"="Or "&"]="And "~"]="Exor

```

Each individual is composed of simple instructions (program lines) between variables and input and output parameters. The input is in the form of sensor values, and is represented as register variables (s_i). The resulting output or action in the form of propulsion control parameters or motor speeds is also given as registers (motor1 and motor2). The same basic evolutionary algorithm is at the heart of both a first system (the non-memory system) which learns directly from actual sensor input, and a second system (the memory system) which learns indirectly from stored experiences.

In the preferred embodiment of the memory system, the population of programs is initialized with random content at the outset. A simple tournament is used for the competitive selection of individuals to produce the right distribution of offspring. The GP-system with its simple steady state tournament selection algorithm has the following execution cycle:

1. Select four arbitrary programs from the population.
- 25 2. For each of the programs calculate fitness.
3. Make two copies (offspring) of the two individuals with highest fitness and let the copies be subject to crossover and mutation.
4. Replace the two individuals with worst fitness with the two new offspring.
5. Repeat step 1 to 4.

30 Symbolic Regression

Both of the preferred embodiments of the present invention, the memory system and the non-memory system, use a fitting algorithm. For both embodiments, the preferred fitting algorithm is a symbolic regression algorithm. Symbolic regression is the procedure of inducing a symbolic equation, function or program which fits given numerical data. Genetic programming is ideal for symbolic regression, and most GP applications could be reformulated as a variant of symbolic regression. A GP system performing symbolic regression takes a number of numerical input/output relations, called fitness cases, and produces a function or machine language computer

program that is consistent with these fitness cases. Consider for example the following fitness cases:

$$f(2) = 6$$

$$f(4) = 20$$

$$f(5) = 30$$

$$5 \quad f(7) = 56$$

These input/output pairs or fitness cases are consistent with the function below:

$$f(x) = x*x+x$$

This very simple example would, in our register machine language, look as:

$$a = x*x;$$

$$10 \quad y_out = a+x;$$

or:

$$a=x+1;$$

$$y_out = x*a;$$

Here the input and the expected output both consist of a single number, but in many cases symbolic
 15 regression is performed with vectors specifying the input/output relation of the sought function. In the examples below the input vector has more than 10 components, and the output vector has in some cases two outputs. The fitness used to guide the system during evolution is often some kind of error summation of the expected values versus the actual values produced by an individual program.

Khepera Robot Experiments

20 Experiments for both the memory and non-memory embodiments of the present invention were performed with a standard autonomous miniature robot, the Swiss mobile robot platform Khepera, which is illustrated in FIG. 23 and designated by the reference numeral 100. It is equipped with eight infrared proximity sensors. The mobile robot has a circular shape, a diameter of 6 cm and a height of 5 cm. It possesses two motors and an on-board power supply. The motors can be independently controlled by a PID controller. The eight
 25 infrared sensors are distributed around the robot in a circular pattern. They emit infrared light, receive the reflected light and measure distances in a short range: 2-5 cm. The robot is also equipped with a Motorola 68331 micro-controller which can be connected to a workstation via a serial cable.

It is possible to control the robot in two ways. The controlling algorithm can be run on a workstation, with data and commands communicated through the serial line. Alternatively, the controlling algorithm can be
 30 cross-compiled on the workstation and down-loaded to the robot which then runs the complete system in a stand-alone fashion. At present, two versions of the system have been developed; one where the controlling GP-system is run on the workstation, and another where the system is downloaded and run autonomously on the micro-controller of the robot.

The micro-controller has 250KB of RAM and a large ROM containing a small operating system. The
 35 operating system has simple multi-tasking capabilities and manages the communication with the host computer. The robot has several extension ports where peripherals such as grippers and TV cameras can be attached.

Training Environment

The training environment used for the experimental obstacle avoiding task used in experiments on both the non-memory and memory embodiments of the present invention was about 70 cm x 90 cm. It has an irregular boarder with different angles and four deceptive dead-ends in each corner. In the large open area in the middle, movable obstacles can be placed. The friction between wheels and surface was low, enabling the robot to slip with its wheels during a collision with an obstacle. There is an increase in friction with the walls making it hard for the circular robot to turn while in contact with a wall.

Objectives

The goal of the GP systems in the experiments was to evolve obstacle avoiding behavior in a sense-think-act context. Both the memory and non-memory systems operate in real-time and aim at obstacle avoiding behavior derived from real noisy sensorial data. The sensorial data for both systems come from eight proximity sensors mounted on the circular platform.

Symbolic regression is, as mentioned above, the procedure of inducing a symbolic equation fitting given numerical data. Genetic programming is ideal for symbolic regression, and most GP applications could be reformulated as variants of symbolic regression. In the non-memory embodiment of the present invention, which involves obstacle avoidance, we approximate a function that takes the sensor values as an input vector and returns an action in the form of a vector of two motor speeds:

$$f(s_1, s_2, s_3, s_4, s_5, s_6, s_7, s_8) = \{m_1, m_2\} \quad (\text{equ. 1})$$

This vector models the simple stimulus-response behavior of the robot. The first system (the non-memory embodiment of the present invention) evolves the function directly through interaction with the environment. The second approach (the memory embodiment of the present invention) evolves a simulation or world model which defines a relationship between inputs (sensor values), outputs (motor speeds) and corresponding predicted fitness values as follows:

$$f(s_1, s_2, s_3, s_4, s_5, s_6, s_7, s_8, m_1, m_2) = \text{predicted fitness} \quad (\text{equ. 2})$$

The second embodiment of the invention is memory-based in that a sensory-motor (input-output) state is "associated" with a fitness, that might be termed "feeling".

Fitness Calculation

The preferred fitness function for the present embodiments of the present invention is an empirically derived fitness function defining the obstacle avoiding task has a pain and a pleasure part. The negative contribution to fitness, called pain, is simply the sum of all proximity sensor values. The closer the robot's sensors are to an object, the more pain it experiences. In order to keep the robot from standing still or gyrating, it has a positive contribution to fitness, called pleasure, as well. It receives pleasure from going straight and fast. Both motor speed values minus the absolute value of their difference is thus added to the fitness.

Let s_i be the values of the proximity sensors ranging from 0 - 1023, where a higher value means

closer proximity to an object. Let m_1 and m_2 be the left and right motor speeds resulting from an execution of an individual. The values of m_1 and m_2 are in the range of zero to 15. The empirical fitness value can then be expressed more formally as:

$$f = \sum s_i + |8 - m_1| + |8 - m_2| + |m_1 - m_2| \quad (\text{equ. 3})$$

5 The Non-Memory Method

The first method of the invention evolves the controlling function (equ. 1) directly, and fitness is calculated from the current event. In the preferred embodiment of the non-memory method, the fitting algorithm used is symbolic regression using CGPS and the evolved programs are true functions; no side-effects are allowed.

- 10 The learning algorithm had a small population size, typically less than 50 individuals. The individuals use the eight values from the sensors as inputs and produce two output values which are transmitted to the robot as motor speeds. Each individual program did this manipulation independent of the others, and thus stood for an individual behavior of the robot when it was invoked to control the motors. Table 1 gives a summary of the problem and its parameters.
- 15 The modules of the learning system and the execution cycle of the GP system are illustrated in FIGs. 23 and 24 respectively. In both methods, each individual is tested against a different real-time fitness case. This could result in "unfair" comparison where individuals have to maneuver in situations with very different possible outcomes. However, experiments show that over time this probabilistic sampling will even out the random effects in learning, and a set of good solutions survive.

Table:	
Objective:	Obstacle avoiding behavior in real-time
Terminal set:	Integers in the range 0-8192
Function set:	ADD, SUB, MUL, SHL, SHR, XOR, OR, AND
5 Raw and standardized fitness:	Pleasure subtracted from pain value desired value
Wrapper:	None
Parameters:	
Maximum population size:	50
Crossover Prob:	90%
10 Mutation Prob:	5%
Selection:	Tournament Selection
Termination criteria:	None
Maximum number of generations:	None
Maximum number of nodes:	256 (1024)

15

TABLE 1

Results Using Non-Memory Method

Using the non-memory method, the robot shows exploratory behavior from the first moment. This is a result of the diversity in behavior residing in the first generation of programs which has been generated randomly. Naturally, the behavior is erratic at the outset of a run.

20

During the first minutes, the robot keeps colliding with different objects, but as time goes on the collisions become more and more infrequent. The first intelligent behavior usually emerging is some kind of backing up after a collision. Then the robot gradually learns to steer away in an increasingly more sophisticated manner.

After about 40-60 minutes, or 120-180 generation equivalents, the robot has learned to avoid obstacles in the rectangular environment almost completely. It has learned to associate the values from the sensors with their respective location on the robot and to send correct motor commands. In this way the robot is able, for instance, to back out of a corner or turn away from an obstacle at its side. Tendencies toward
 5 adoption of a special path in order to avoid as many obstacles as possible can also be observed. The number of collisions per minute diminishes as the robot learns and the population becomes dominated by good control strategies.

Despite only processing results for less than 0.1% of the CPU time, the method competes well with other paradigms. It is for instance 100 times faster than a related evolutionary approach evolving Neural
 10 Network controllers on the same robot platform.

The moving robot gives the impression of displaying a very complex behavior. Its behavior resembles that of a bug or an ant exploring an environment, with small and irregular moves around the objects.

Memory Based Method

15 The memory-based embodiment of the present invention consists of two separate processes or units. A computing or planning unit 101 communicates with inputs (sensors and motors) as well as storing events in a memory buffer. A computer model unit 103 is continuously learning and inducing or evolving a model of the world consistent with the entries in a memory buffer.

The former process is called the planning process because it is involved in deciding what action to
 20 perform given a certain model of the world. The latter process is called the learning process because it consists of trying to derive a model (in the form of a function) from memory data.

The present invention, in its embodiment as a memory based control system as illustrated in FIG.
 25 includes six major components.

1. The robot 100 with sensors and actuators.
- 25 2. A memory buffer 102 which stores event vectors representing events in the past.
3. An evolution unit in the form of a GP system 104 which evolves a model of the world that fits the information of the event vectors.
4. A fitness calculation module 106 which calculates an empirical fitness.
5. A currently best induced individual computer model 108.
- 30 6. A search module 110 that determines the best action given the currently best world model.

The Planning Process

The main execution cycle of the planning process is illustrated in FIG. 27, and has several similarities with the execution cycle of the simple genetic control architecture. The planning unit 101 has
 35 actual communication with the robot 100 and decides what action should be performed next. It accesses the

best model 108 of the world supplied by the learning unit 103. The planning unit 101 has three main objectives. It communicates with the robot 100, finds a feasible action, and stores the resulting event in the memory buffer 102.

As illustrated in FIG. 27, the process starts with reading all eight infrared proximity sensors.

5 These values are used to instantiate the corresponding variables in the currently best world model 108. The next objective is to find a favorable action given the current sensor values. In the obstacle avoiding task used here the possible actions are 16 different motor speeds for each of the two motors. Each motor has 8 forward speeds, 7 backward speeds, and a zero speed. Combining all alternatives of the two motors, there are 256 different actions altogether to choose from.

10 This comparatively small figure means that we can easily afford to search through all possible actions while the world model 108 provides a predicted fitness for each of them. The induced model in the form of a computer program 108 from the learning unit 103 can thus be seen as a simulation of the environment consistent with past experiences, where the robot 100 can simulate different actions. The action which gives the best fitness is remembered and sent as motor speeds to the robot 100.

15 If an autonomous agent (e.g. robot 100) had so many possible atomic actions that an exhaustive search would be infeasible, we could use a heuristic search method to find an action that gives good predicted fitness. We would, for instance, have another genetic programming system evolving a stepwise plan which optimizes fitness according to the currently best world model.

In order to get feed-back from the environment, the planning unit 101 has to sleep and await the
20 result of the chosen action. The planning unit 101 sleeps 300 ms while the robot 100 performs the movement defined by the motors speeds. This time is an approximate minimum in order to get usable feed-back from changes in the sensor values in the present example. Thus, the main operation of the planning unit 103 is the sleeping period waiting for feedback from the environment and it, therefore, consumes less than 0.1% of the total CPU time of the system.

25 After the sleeping period the sensor values are read again. These new values are used to compute a new empirical fitness value using (equ. 3). This fitness value is stored, together with the earlier sensor values and the motor speeds as an event vector. The event vector consists of 11 numbers; the eight sensor values, the two motor speeds, and the resulting calculated (empirical) fitness. This vector represents what the agent experienced, what it did, and what the results were of its action. The memory buffer stores 50 of these
30 events and shifts out old memories according to a predetermined schema.

It is then the responsibility of the learning unit 103 to evolve a program that simulates the environment as good and as consistent as possible with respect to the events in the memory buffer 102. As will be seen below, this can be done by a straight-forward application of symbolic regression through genetic programming.

The Learning Process

The objective of the learning unit 103 is to find a function or a program which will calculate the predicted fitness of an action (set of motor speed outputs), given the initial conditions in form of the sensor values:

$$f(s_1, s_2, s_3, s_4, s_5, s_6, s_7, s_8, m_1, m_2) = \text{predicted fitness} \quad (\text{equ. 4})$$

In most cases an additional state variable is used as part of the program. This is simply a memory cell allowing the possibility to use side-effects in evolved solutions. This feature significantly improves performance in the complex environment used throughout these experiments:

$$f(s_1, s_2, s_3, s_4, s_5, s_6, s_7, s_8, m_1, m_2, \text{state}_i) = \{\text{predicted fitness}, \text{state}_i + 1\} \quad (\text{equ. 5})$$

- 10 Each event vector stores an instance of the values: the inputs (sensor values), outputs (motor speeds) and the resulting empirical fitness.

FIG. 26 illustrates the interactions between the GP system 104 and the memory buffer 102 in the learning process.

Giving the System a Childhood

- 15 A simple approach to managing the memory buffer 102 when all 50 places have been filled is to shift out the oldest memories as the new entries come in. However, it was determined that the system then forgot important early experiences. Early mistakes, made before a good strategy was found, are crucial to remember in order to not evolve world models that permit the same mistakes to be made again. Therefore, the robot is given a childhood - an initial period whose memories were not so easily forgotten. The childhood
20 also reduces the likelihood that the system displays a strategy sometimes seen even in humans - it would only perform actions which confirmed its current (limited) world model.

- Another important factor for successfully inducing an efficient world model is to have a stimulating childhood. It is important to have a wide set of experiences to draw conclusions from. Noise is therefore added to the behavior in the childhood to avoid stereotypic behavior very early in the first seconds of
25 the system's execution. As long as experiences are too few to allow for a meaningful model of the world, this noise is needed to assure enough diversity in early experiences. The childhood of the system is defined as the time before the memory buffer is filled which takes about 20 seconds.

Results

- The memory based system quickly learns the obstacle avoiding task in most individual
30 experiments. It normally takes only a few minutes before the robot displays a successful obstacle avoiding behavior. The obvious reason for the speed up using memory can be identified in the flowchart of the algorithms.

- In the second (memory) method, there is no "sleeping" period which means that the genetic programming system can run at the full speed possible for the CPU. This results in a speed up of more than
35 2000 times in the GP system. On the other hand, there is now a more complex task to learn. Instead of

evolving an ad-hoc strategy for steering the robot, the system now has to evolve a complete model of relationships between the eight input variables, the two action variables and the fitness.

The population size is preferably increased from 50 individuals to 10,000 to ensure robust learning in the memory-based system as illustrated in TABLE 2. The system still has to wait for the robot 5 100 to collect enough memory events to draw some meaningful conclusions. Yet the speed-up with memory exceeds a factor of 40 which makes it possible for the system to learn a successful behavior in less 1.5 minutes on average. This means that the behavior emerges 4000 times faster than in similar approaches.

	Objective:	Symbolic regression of environment simulation
	Terminal set:	Integers in the range 0-8192
10	Function set:	ADD, SUB, MUL, SHL, SHR, XOR, OR, AND
	Raw and standardized fitness:	The sum taken over 50 fitness cases of the absolute value of difference between the actual and desired value predicting the fitness.
	Wrapper:	None
	Parameters:	
	Maximum population size:	10000
15	Crossover Prob:	90%
	Mutation Prob:	5%
	Selection:	Tournament Selection
	Termination criteria:	None
	Maximum number of generations:	None
20	Maximum number of nodes:	256 (1024)

TABLE 2

The behavior of the robot is very different between the two systems (memory and non-memory) discussed here. The system without memory behaves in a very complex way and gives the impression of a small bug which randomly runs around avoiding obstacles, but with little overall strategy.

25 The memory system, on the other hand, displays a set of very "thought through" behaviors. The robot always displays a clear strategy and travels in straight lines or smooth curves. Some of the behaviors evolved show an almost perfect solution to the current task and fitness function.

The robot usually demonstrates a limited set of strategies during evolution in our experiments.

Some of the emerging intelligent strategies might be described as belonging to different behavioral classes (ordered according to increasing success):

1. The straight and fast strategy: This is the simplest "intelligent" behavior. The induction process has only seen the pattern arising from the pleasure part of the fitness function. The model of the robot
5 and its environment thus only contains the relationship expressing that going straight and fast is good. The robot consequently heads into the nearest wall and continues to stand there spinning its wheels. This strategy sometimes emerges right after the childhood when the noise is removed and the system is solely controlled by inferences from the induced model.
2. The turning behavior: The second simplest strategy is based on the experience that
10 turning often improves fitness. The robot starts spinning around its own axis and does avoid all obstacles but also ignores the pleasure part of the fitness rewarding it for going straight and fast.
3. The dancing strategy: This strategy uses the state information in the model and navigates to the open space where it starts to move in an irregular circular path avoiding obstacles. Most of the time the robot moves around keeping a distance to the obstacles big enough to avoid any reaction from its sensors. If
15 this strategy worked in all cases it would be nearly perfect because it keeps obstacles out of reach of the sensors and the robot is totally unexposed to pain. In most cases, however, the robot wanders off its path and comes too close to an obstacle where it then is unable to cope with the new situation and experiences collisions.
4. The backing-up strategy: This is the first effective technique that allows the robot to
20 avoid obstacles while moving around to some extent. The path the robot travels is, however, very limited and it is not a preferred solution.
5. The bouncing strategy. Here the robot gradually turns away from an obstacle as it approaches it. It looks as if the robot bounces like a ball at something invisible close to the obstacle. This behavior gives a minimum speed change in the robot's path.
- 25 6. The perfect or nearly perfect strategy: The robot uses the large free space in the middle of the training environment to go straight and fast, optimizing the pleasure part of the fitness. As soon as the robot senses an object it quickly turns 180 degrees on the spot and continues going straight and fast. This strategy also involves state information because turning 180 degrees takes several events in the robot's perception, and that cannot be achieved without states.

30 Most of the experiments displayed a simple behavior very early after the childhood, just to realize a more successful pattern a few seconds later and change its strategy correspondingly. The change in strategy was always accomplished by a new best individual and fitness value displayed by the GP algorithm.

TABLE 3 illustrates the results of 10 evaluation experiments with the memory based system. The results were produced by timing the robot's behavior in 10 consecutive experiments. In each experiment the
35 robot was watched for 20 minutes before the experiment was terminated. Each time the behavior changed was noted. The table gives the number of the experiment, the strategy displayed when the experiment was terminated and the time when this strategy first appeared.

It is not completely evident what really constitutes an autonomous agent. Some would argue that the autonomy is a property of the controlling algorithm while others would argue that physical autonomy is needed.

In order to try the fully autonomous paradigm, a special version of the system was ported to the 5 micro-controller.

Run number	Result	Time (minutes)
1	perfect	1.5
2	perfect	1.5
3	perfect	0.5
10 4	perfect	3.0
5	perfect	2.0
6	perfect	2.0
7	perfect	1.5
8	perfect	2.0
15 9	perfect	1.0
10	dancing	0.5

TABLE 3

It is possible to download these systems via the serial cable to the robot. With the accumulators switched on, the robot can then be disconnected from the workstation and can run fully autonomous. The 20 Motorola 68331 micro-controller then runs the complete GP learning system.

As mentioned earlier, the micro controller has 256 KB of RAM memory. The kernel of the GP system occupies 32 KB, and each individual 1KB, in the experimental setup. The complete system without memory consists of 50 individuals and occupies 82KB which is well within the limits of the on-board system. The more complex system, learning from memory, has to use a smaller population size than the desired 10000 25 individuals. This results in less robust behavior with a more frequent convergence to local optima such as displayed by the first strategies in Figure 15.

In either case, it is demonstrated herewith that the compactness of the compiling GP system enables relatively powerful solutions in weak architectures such as those used in embedded control.

Summary of Autonomous Agent Control System Experimental Results

It has been demonstrated that a GP system can be used to control an existing robot in a real-time environment with noisy input. The evolved algorithm shows robust performance even if the robot is lifted and placed in a completely different environment or if obstacles are moved around. It is believed that the robust behavior of the robot partly could be attributed to the built-in generalization capabilities of the genetic programming system.

We have also cross-compiled the GP-system and run it in the same set-up on the micro-controller on board the robot. This demonstrates the applicability of Genetic Programming to control tasks on low-end architectures. The technique could potentially be applied to many one-chip control applications in, for instance, consumer electronics devices.

It has been demonstrated that the use of memory with a real robot and a GP based control system could speed up learning by a factor of 40. Furthermore, the strategies evolved with the memory based system display a smoother, less chaotic behavior, undisturbed by the internal dynamics of the GP system. The memory-based system could also be expected to scale up better because training times are not directly related to the dynamics of the agent and its environment, but instead almost completely depend on the difficulty of the induction problem in the application domain.

(a)

In summary, the present invention overcomes the drawbacks of the prior art by eliminating all compiling, interpreting or other steps that are required to convert a high level programming language instruction such as a LISP S-expression into machine code prior to execution or that are required to access Learned Elements or run-time data in data structures.

This makes possible the practical solutions to problems which could not heretofore be solved due to excessive computation time. For example, a solution to a difficult problem can be produced by the present system in hours, whereas a comparable solution might take years using conventional techniques.

Various modifications will become possible for those skilled in the art after receiving the teachings of the present disclosure without departing from the scope thereof.

INDUSTRIAL APPLICABILITY

The present invention has utility in computerized learning which can be used to generate solutions to problems in numerous technical areas, and also to control of an autonomous agent such as an industrial robot.

WE CLAIM:

1. A computer implemented learning method characterized by the steps of:
 - (a) providing a computer with an indirectly executable program including:
 - a first instruction that points to and designates machine code stored in a
5 memory as data;
 - a second instruction that points to and designates machine code stored in a
memory as directly executable code;
 - a third instruction that alters machine code pointed to by the first instruction;
and
 - 10 a fourth instruction that executes machine code pointed to by the second
instruction; and
 - (b) controlling the computer to execute the program which performs the steps of:
 - (b1) creating and storing a machine code entity including a directly executable
instruction in a memory;
 - 15 (b2) executing the second instruction to point to the entity;
 - (b3) executing the fourth instruction using input data to produce a result;
 - (b4) evaluating the result;
 - (b5) executing the first instruction to point to the entity;
 - (b6) executing the third instruction to alter the entity to include a different
20 directly executable instruction; and
 - (b7) repeating steps (b3) to (b6) until an end criterion is reached.
2. A method as in claim 1, in which:
 - the first instruction points to and designates machine code stored in a memory as data by
casting a pointer thereto as a data pointer; and
 - 25 the second instruction points to and designates machine code stored in a memory as a
directly executable code by casting a pointer thereto as a function pointer.
3. A method as in claim 1, in which:
 - the first instruction points to and designates machine code stored in a memory as data by
creating a data pointer thereto; and
 - 30 the second instruction points to and designates machine code stored in a memory as
directly executable code by creating a function pointer thereto.
4. A method as in claim 1, in which:
 - step (b1) comprises creating the entity as comprising a header and a footer; and
 - step (b6) comprises preventing the header and footer from being altered.
- 35 5. A method as in claim 1, in which:
 - step (b1) comprises creating the entity as comprising a return instruction; and

- step (b6) comprises preventing the return instruction from being altered.
6. A method as in claim 1, in which the learning algorithm in step (b4) and the alteration algorithm in step (b6) in combination comprise a genetic algorithm.
7. A method as in claim 1, in which step (b1) comprises creating the entity as comprising a
5 plurality of directly executable instructions.
8. A method as in claim 7, in which step (b1) comprises randomly selecting the directly executable instructions from a set of directly executable instructions.
9. A method as in claim 7, in which step (b6) comprises genetically altering the directly executable instructions.
10. A method as in claim 7, in which:
one of the directly executable instructions comprises an operator and an operand; and
step (b6) comprises genetically mutating the operand.
11. A method as in claim 7, in which two of the directly executable instructions comprise operators respectively; and
15 step (b6) comprises genetically crossing the operators.
12. A method as in claim 1, in which:
step (b1) comprises randomly selecting the directly executable instruction from a set of directly executable instructions.
13. A method as in claim 1, in which step (b6) comprises genetically altering the directly
20 executable instruction.
14. A method as in claim 1, in which:
the directly executable instruction comprises an operator and an operand; and
step (b6) comprises genetically mutating the operand.
15. A computer implemented learning method characterized by the steps of:
25 (a) providing a computer with an indirectly executable program; and
(b) controlling the computer to execute the program which performs the steps of:
(b1) creating a directly executable machine code entity including a directly
executable program instruction and a directly executable return instruction;
(b2) executing the entity using input data to produce a result;
30 (b3) evaluating the result;
(b4) altering only the program instruction to include a different directly
executable program instruction; and
(b6) repeating steps (b2) to (b4) until an end criterion is reached.
16. A method as in claim 15, in which step (b4) comprises preventing the return instruction
35 from being altered.
17. A method as in claim 15, in which step (b4) comprises allowing the return instruction to be altered, and then restoring the return instruction to its original state.

18. A method as in claim 15, in which step (b4) comprises separating the program instruction from the entity, altering the program instruction to include the different program instruction, and returning the different program instruction to the entity.

19. A method of manipulating and executing computer machine code, comprising the steps
5 of:

(a) providing a computer with an indirectly executable program including:

a first instruction that points to and designates machine code stored in a memory
as data;

a second instruction that points to and designates machine code stored in a
10 memory as directly executable code;

a third instruction that alters machine code pointed to by the first instruction;
and

a fourth instruction that executes machine code pointed to by the second
instruction; and

15 (b) controlling the computer to execute the program which performs the steps of:

(b1) creating and storing a machine code entity including a directly executable
instruction in a memory;

(b2) executing the second instruction to point to the entity;

(b3) executing the fourth instruction to execute the entity;

20 (b4) executing the first instruction to point to the entity; and

(b5) executing the third instruction to alter the entity to include a different
directly executable instruction.

20. A method as in claim 19, in which:

step (b1) comprises creating the entity as including a directly executable program
25 instruction and a directly executable return instruction; and

step (b5) comprises preventing the return instruction from being altered.

21. A computer learning system characterized by comprising:

a memory for storing an indirectly executable computer program including:

a first instruction that points to and designates machine code stored in the
30 memory as data;

a second instruction that points to and designates machine code stored in the
memory as directly executable machine code;

a third instruction that alters machine code pointed to by the first instruction;
and

35 a fourth instruction that executes machine code pointed to by the second
instruction; and

a processor for executing the program;

the processor, memory and program operating in combination for performing the steps of:

- (a) creating and storing a machine code entity including a directly executable instruction in the memory;
- 5 (b) executing the second instruction to point to the entity;
- (c) executing the instruction using input data to produce a result;
- (d) evaluating the result;
- (e) executing the first instruction to point to the entity;
- (f) executing the third instruction to alter the entity using an alteration algorithm
- 10 to include a different directly executable instruction; and
- (g) repeating steps (b) to (f) until an end criterion is reached.

22. A system as in claim 21, in which:

- the first instruction points to and designates machine code stored in the memory as data by casting a pointer thereto as a data pointer; and
- 15 the second instruction points to and designates machine code stored in the memory as directly executable code by casting a pointer thereto as a function pointer.

23. A system as in claim 21, in which:

- the first instruction points to and designates machine code stored in the memory as data by creating a data pointer thereto; and
- 20 the second instruction points to and designates machine code stored in the memory as directly executable code by creating a function pointer thereto.

24. A system as in claim 21, in which:

- step (a) comprises creating the entity as comprising a header and a footer; and
- step (f) comprises preventing the header and footer from being altered.

25. A system as in claim 21, in which:

- step (a) comprises creating the entity as comprising a return instruction; and
- step (f) comprises preventing the return instruction from being altered.

26. A system as in claim 21, in which the learning algorithm in step (d) and the alteration algorithm in step (f) in combination comprise a genetic algorithm.

30 27. A system as in claim 21, in which step (a) comprises creating the entity as comprising a plurality of directly executable instructions.

28. A system as in claim 27, in which step (a) comprises randomly selecting the directly executable instructions from a set of directly executable instructions.

29. A system as in claim 27, in which step (f) comprises genetically altering the directly

35 executable instructions.

30. A system as in claim 27, in which:

- one of the directly executable instructions comprises an operator and an operand; and

step (f) comprises genetically mutating the operand.

31. A system as in claim 27, in which two of the directly executable instructions comprise operators respectively; and

step (f) comprises genetically crossing the operators.

5 32. A system as in claim 21, in which step (a) comprises randomly selecting the directly executable instruction from a set of directly executable instructions.

33. A system as in claim 21, in which step (f) comprises genetically altering the directly executable instruction.

34. A system as in claim 21, in which:
10 the directly executable instruction comprises an operator and an operand; and
step (f) comprises genetically mutating the operand.

35. A system as in claim 21, in which:
the memory comprises a main memory and a processor memory that is part of the
processor; and

15 step (a) comprises storing the entity in the processor memory.

36. A system as in claim 35, comprising an integrated circuit chip, in which the processor
and the processor memory are formed on the chip.

37. A system as in claim 36, in which at least a portion of the program is stored processor
memory.

20 38. A system as in claim 37, in which:
the processor memory comprises a non-volatile memory section; and
said at least a portion of the program is stored in the non-volatile memory section.

39. A method of creating and executing machine code using a digital computer which
includes a processor and a memory, characterized by the steps of:
25 (a) selecting a directly executable machine code program instruction;
(b) determining a binary number which represents the program instruction to the
processor;
(c) selecting a location in the memory;
(d) creating a directly executable machine code entity which includes the binary number;
30 (e) storing the entity as data in said location;
(f) providing an indirectly executable instruction which causes the processor to directly
execute machine code stored in said location; and

(g) controlling the processor to process the indirectly executable instruction, resulting in
the processor directly executing the entity as including the program instruction.

35 40. A method as in claim 39, in which step (f) comprises providing the indirectly executable
instruction such that said processing thereof in step (g) includes compiling.

41. A method as in claim 39, in which step (f) comprises providing the indirectly executable

instruction such that said processing thereof in step (g) includes interpreting.

42. A method as in claim 39, in which step (f) comprises creating the indirectly executable instruction such that said processing thereof in step (g) includes assembling.

43. A method as in claim 39, further comprising the step of:
5 (h) repeating step (g) at least one time.

44. A method as in claim 39, further comprising the step of:
(h) evaluating a result produced by executing the program instruction in step (g).

45. A method as in claim 39, further comprising the step, performed after step (g), of:
(h) altering the entity as data such that the binary number is changed to represent a
10 different directly executable machine code program instruction to the processor.

46. A method as in claim 45, further comprising the step, performed after step (h), of:
(i) controlling the processor to process the indirectly executable instruction, resulting in the processor directly executing the entity as including the program instruction.

47. A method as in claim 46, further comprising the step of:
15 (j) evaluating a result produced by executing the program instruction in step (i).

48. A method as in claim 39, further comprising the steps, performed after step (g), of:
(h) copying the entity to produce a copied entity; and
(i) altering the copied entity as data such that the binary number is changed to represent a different directly executable machine code program instruction to the processor.

20 49. A method as in claim 39, further comprising the steps of:
(h) evaluating a result produced by executing the program instruction in step (g);
(i) altering the entity as data such that the binary number is changed to represent a different directly executable machine code program instruction to the processor;
(j) controlling the processor to process the indirectly executable instruction, resulting in
25 the processor directly executing the entity as including the program instruction; and
(k) evaluating a result produced by executing the program instruction in step (j).

50. A method as in claim 49, further comprising the step of:
(l) comparing the evaluation of step (k) with the evaluation of step (h).

51. A method as in claim 50, further comprising the step of:
30 (m) selecting a subsequent action in accordance with said comparison of step (l).

52. A method as in claim 39, in which step (d) comprises creating the entity to further include a return instruction.

53. A method as in claim 52, further comprising the steps, performed after step (g), of:
(h) altering the entity as data such that the binary number is changed to represent a
35 different directly executable machine code program instruction to the processor; and

(i) preventing the return instruction from being altered while performing step (h).

54. A method as in claim 52, further comprising the steps, performed after step (g), of:

(h) altering the entity as data such that the binary number is changed to represent a different directly executable machine code program instruction to the processor, and the return instruction is allowed to be altered; and

(i) restoring the return instruction to original state after alteration in step (h).

5 55. A method as in claim 39, in which step (d) comprises creating the entity to further include at least one of a header and a footer.

56. A method as in claim 55, further comprising the steps, performed after step (g), of:

(h) altering the entity as data such that the binary number is changed to represent a different directly executable machine code program instruction to the processor; and

10 (i) preventing said at least one of a header and a footer from being altered while performing step (h).

57. A method as in claim 55, further comprising the steps, performed after step (g), of:

(h) altering the entity as data such that the binary number is changed to represent a different directly executable machine code program instruction to the processor, and said at least one of a
15 header and a footer is allowed to be altered; and

(i) restoring said at least one of a header and a footer to original state after alteration in step (h).

58. A method as in claim 55, further comprising the steps, performed after step (g), of:

(h) separating the binary number from the entity;

20 (i) altering the binary number to produce a different binary number which represents a different directly executable machine code program instruction to the processor; and

(j) replacing the binary number with the different binary number in the entity.

59. A method as in claim 39, in which step (d) comprises creating the program instruction to include run time data.

25 60. A method as in claim 59, further comprising the step, performed after step (g), of:

(h) altering the entity as data such that the run time data is changed.

61. A method as in claim 60, further comprising the step, performed after step (h), of:

(i) controlling the processor to process the indirectly executable instruction, resulting in the processor directly executing the entity as including the program instruction.

30 62. A method as in claim 61, in which step (h) further comprises altering the entity as data such that the binary number is changed to represent a different directly executable machine code program instruction to the processor.

63. A method as in claim 61, further comprising the step of:

(j) evaluating a result produced by executing the program instruction in step (i).

35 64. A method as in claim 59, further comprising the steps, performed after step (g), of:

(h) copying the entity to produce a copied entity; and

(i) altering the copied entity as data such that the run time data is changed.

65. A method as in claim 59, further comprising the steps of:
 (h) evaluating a result produced by executing the program instruction in step (g);
 (i) altering the entity as data such that the run time data is changed;
 (j) controlling the processor to process the indirectly executable instruction, resulting in
 5 the processor directly executing the entity as including the program instruction; and
 (k) evaluating a result produced by executing the program instruction in step (j).
66. A method as in claim 65, further comprising the step of:
 (l) comparing the evaluation of step (k) with the evaluation of step (h).
67. A method as in claim 66, further comprising the step of:
 10 (m) selecting a subsequent action in accordance with said comparison of step (l).
68. A method as in claim 39, in which step (a) comprises selecting the program instruction
 from a set of directly executable program instructions.
69. A method as in claim 39, further comprising the step, performed before step (a), of:
 selecting a subset of program instructions from a set of directly executable program
 15 instructions; and
 step (a) comprises selecting the program instruction from said subset.
70. A method as in claim 39, in which:
 step (d) comprises creating the program instruction to include run time data; and
 the method further comprises the step of:
 20 (h) repeating step (g) for different input values of the run time data.
71. A method as in claim 70, further comprising the step of:
 (i) determining a value of the run time data which produces a predetermined result.
72. A method as in claim 70, in which step (h) further comprises repeating step (g) for
 different program instructions.
- 25 73. A method as in claim 39, in which the method further comprises the step, performed
 after step (g), of:
 (h) creating a representation of the entity in a selected computer programming language.
74. A method as in claim 73, in which step (h) comprises decompiling the entity.
75. A method as in claim 39, in which:
 30 step (d) comprises creating the program instruction to include run time data; and
 the method further comprises the steps of:
 (h) determining if execution of step (g) produced a predetermined result; and
 (i) if step (g) produced said predetermined result, creating a representation of the entity
 including said input value of the run time data in a selected computer programming language.
- 35 76. A method as in claim 75, further comprising the step of:
 (j) repeating step (g) for different input values of the run time data;
 (k) determining a value of the run time data which produced a predetermined result; and

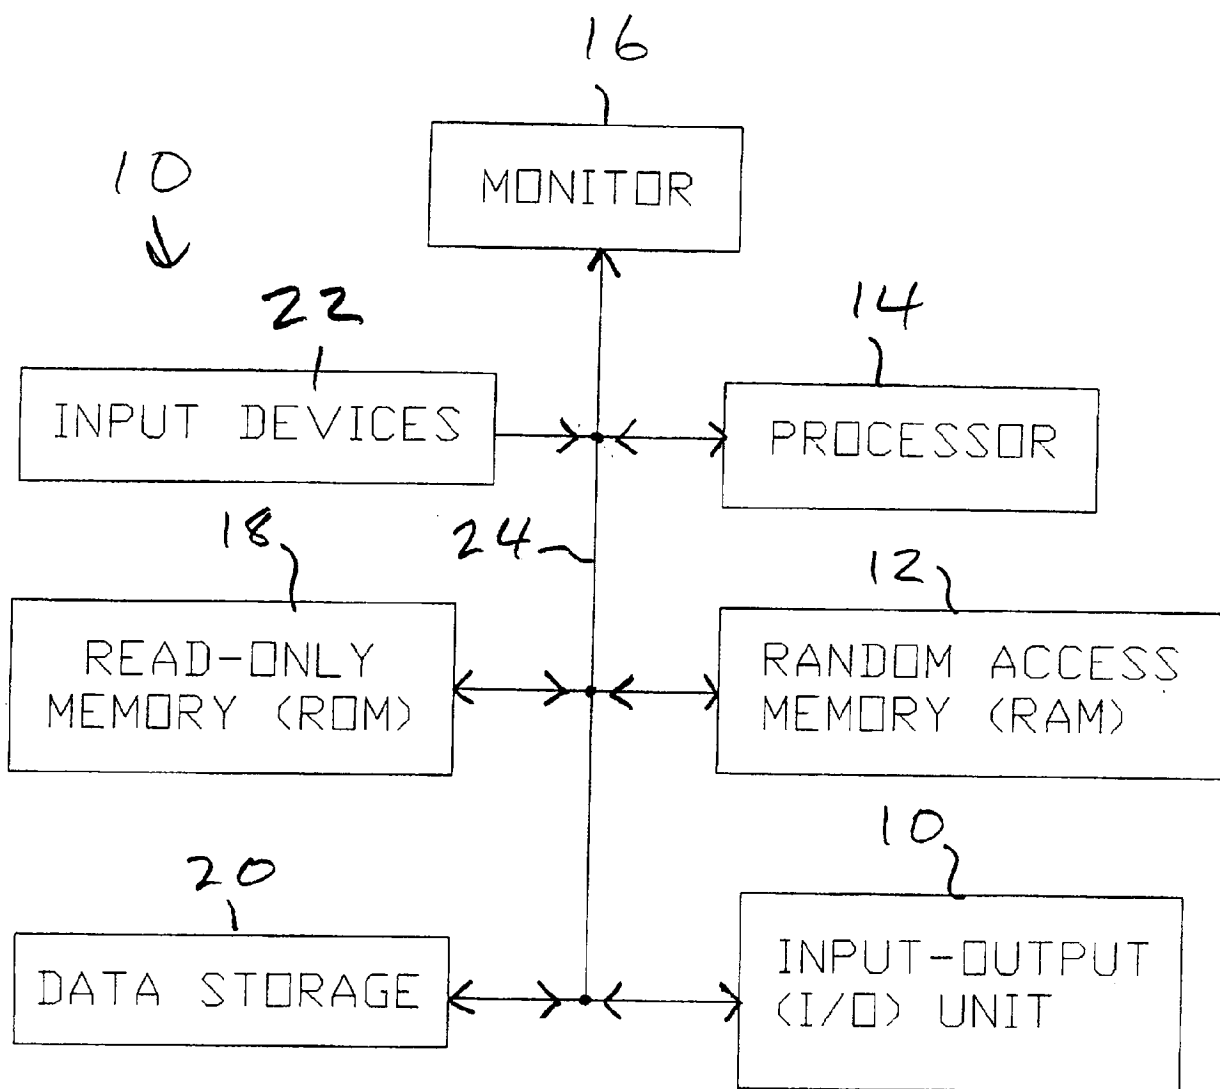


Fig. 1

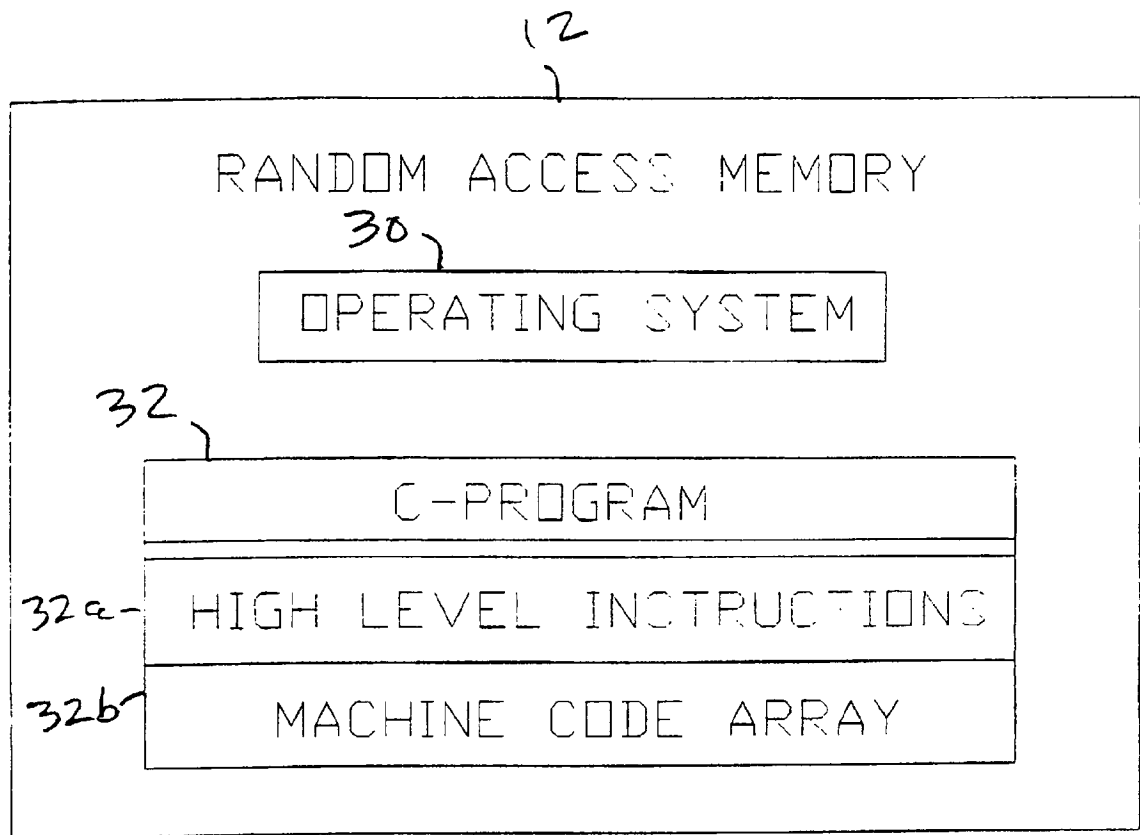


Fig. 2

2/46

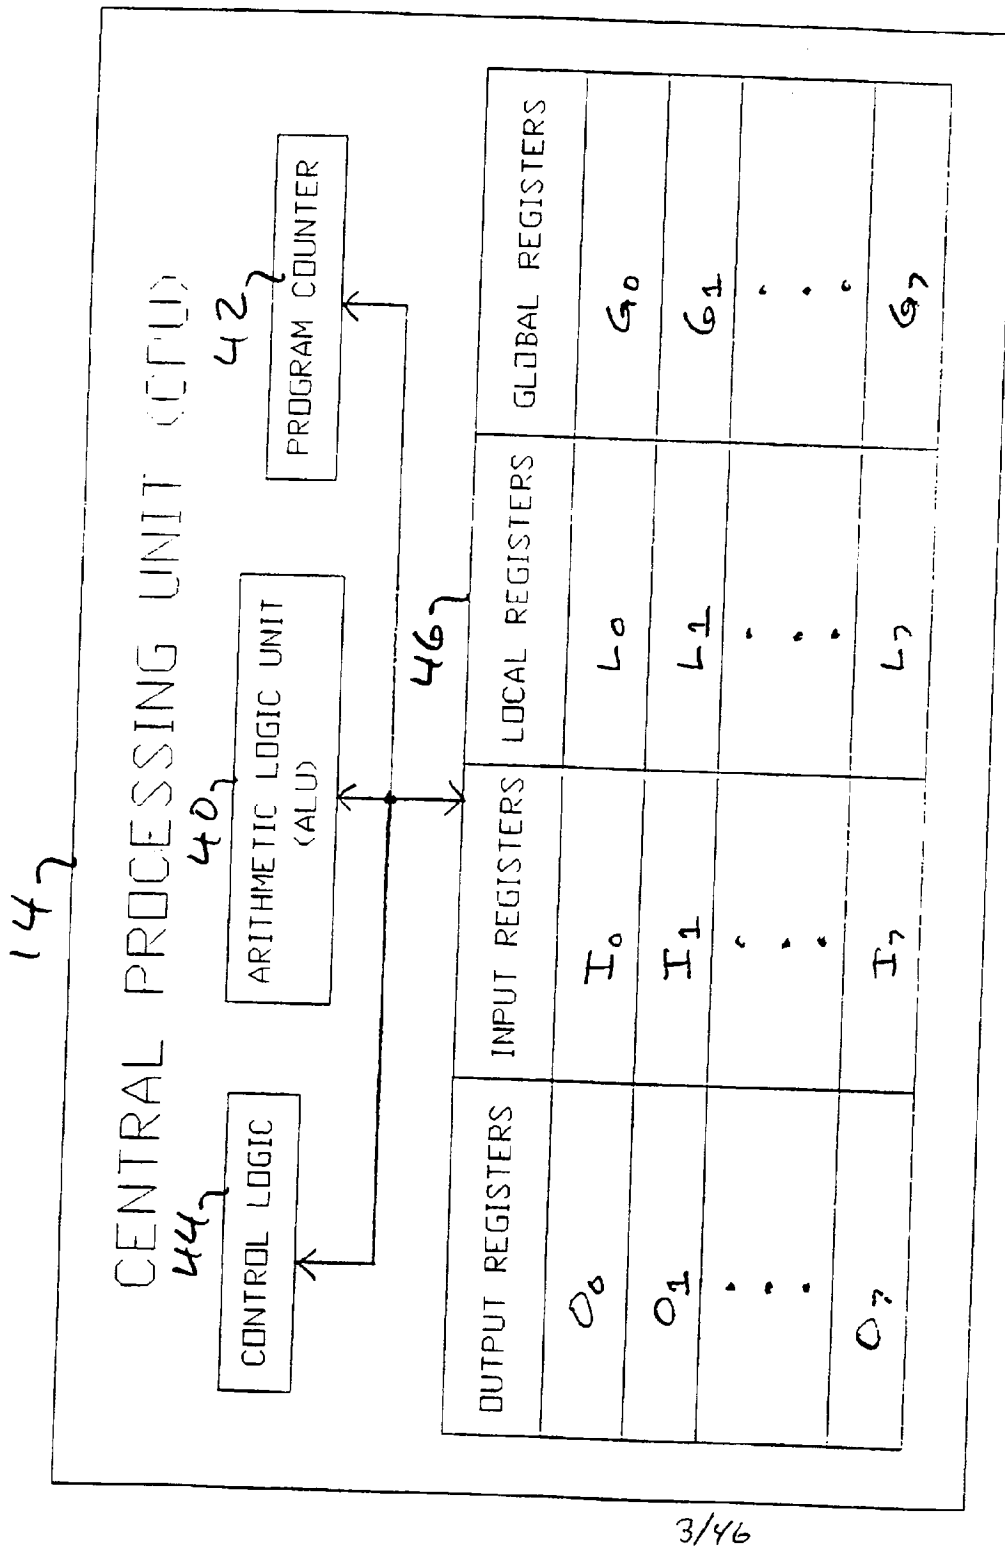


Fig. 3

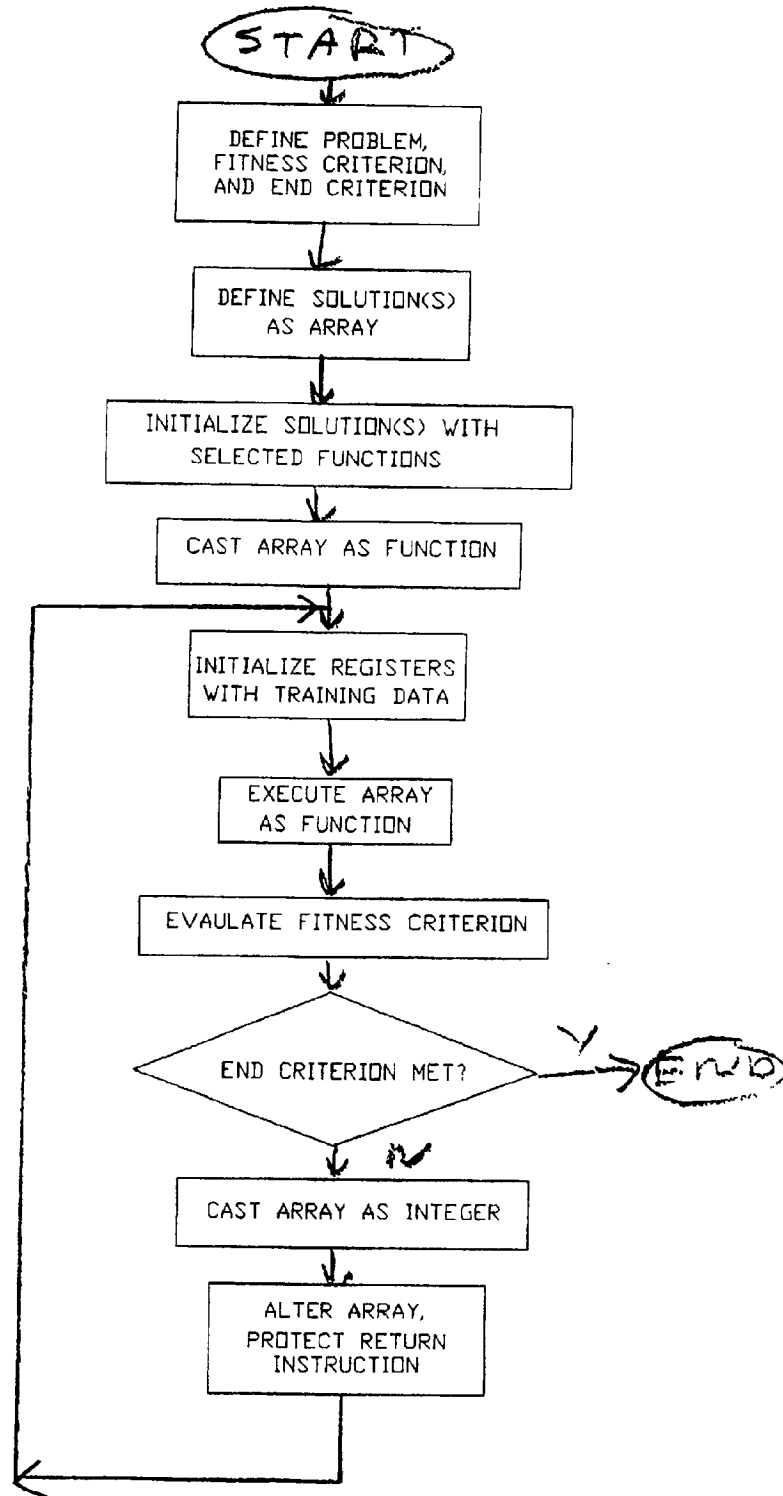


Fig. 4

4/46

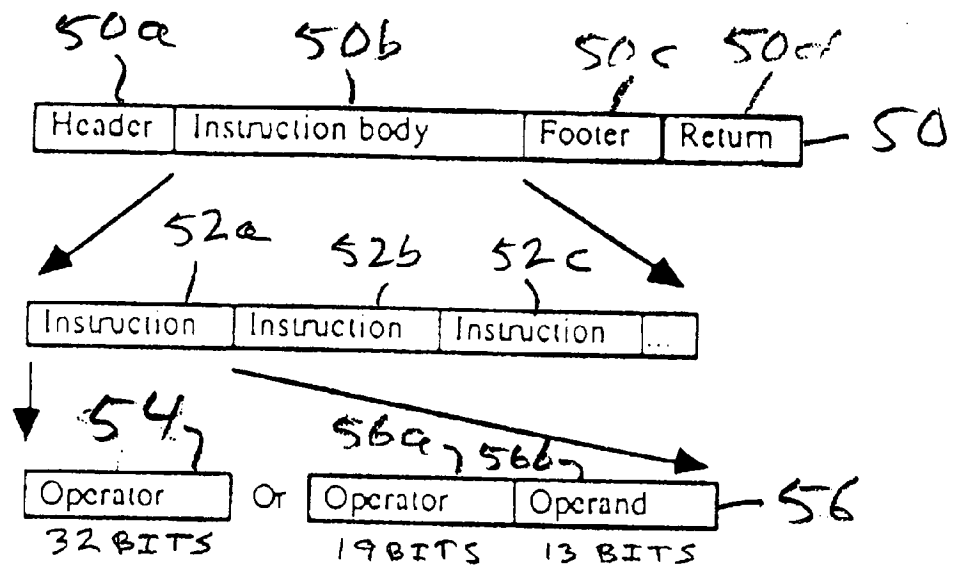


Fig. 5

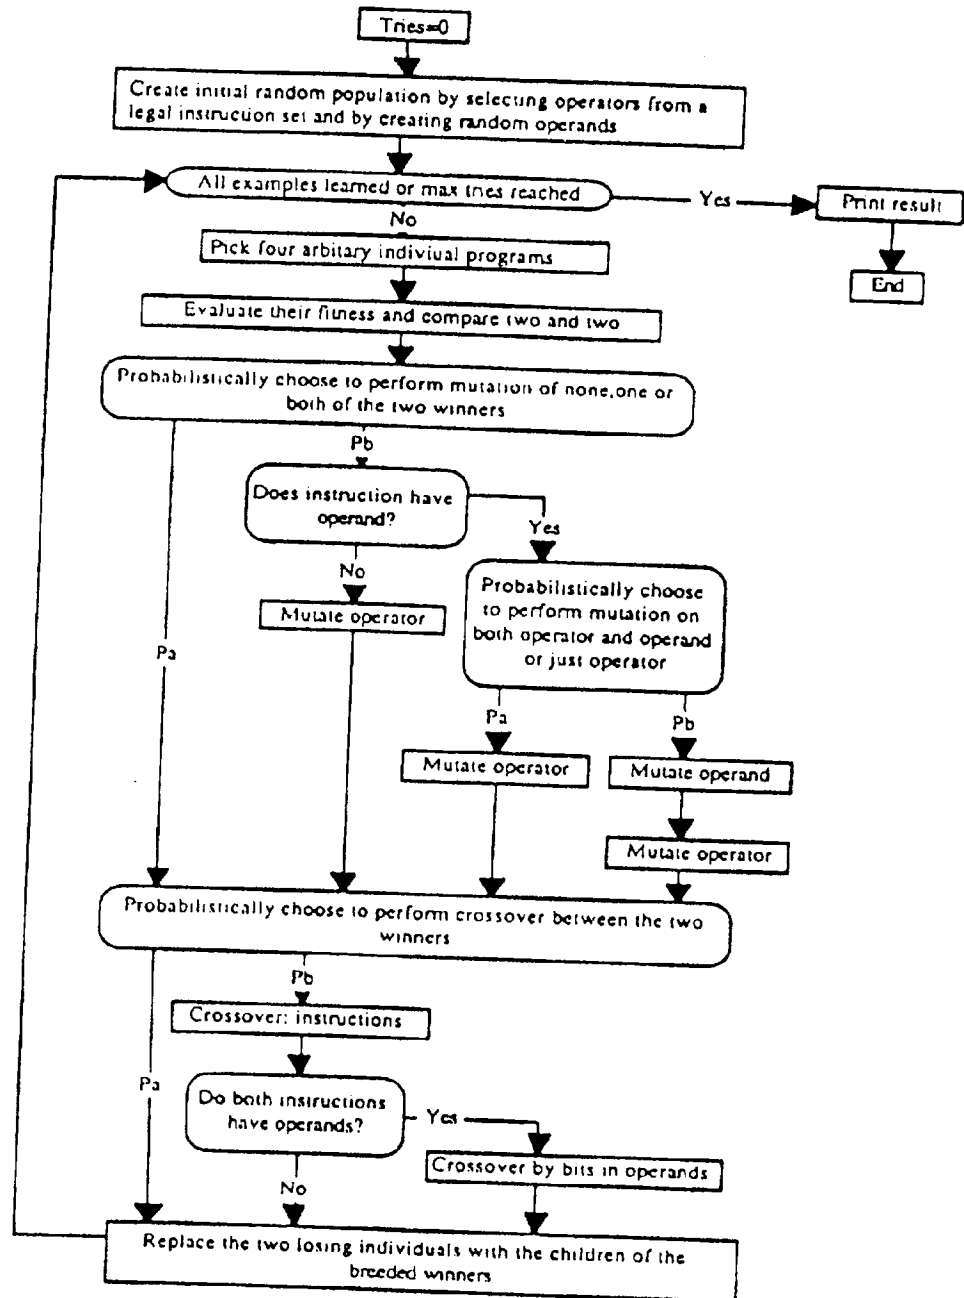


Fig. 6

66
↓





Fig. 7

7/46

Fig. 8






Fig. 9

PARENTS

70-	H	1	3	4	2	7	F	R		
				↕	↕					
72-	H	6	8	5	9	4	8	3	F	R

Fig. 10a

CHILDREN

70'-	H	1	3	5	9	7	F	R
------	---	---	---	---	---	---	---	---

72'-	H	6	8	4	2	4	8	3	F	R
------	---	---	---	---	---	---	---	---	---	---

Fig. 10b

9/46

PARENTS




Fig. 11a

CHILDREN






Fig. 11b

10/46



F i g . 1 2 b





Fig. 13b

12/46




Fig. 14




Fig. 15

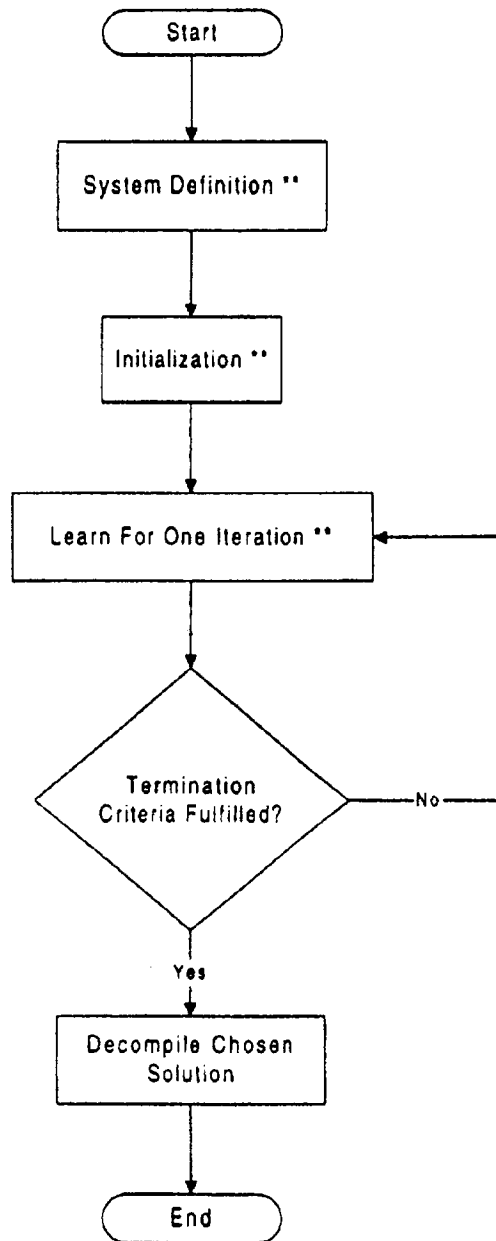


Figure 16a

*General Operation of
the System
14/46*

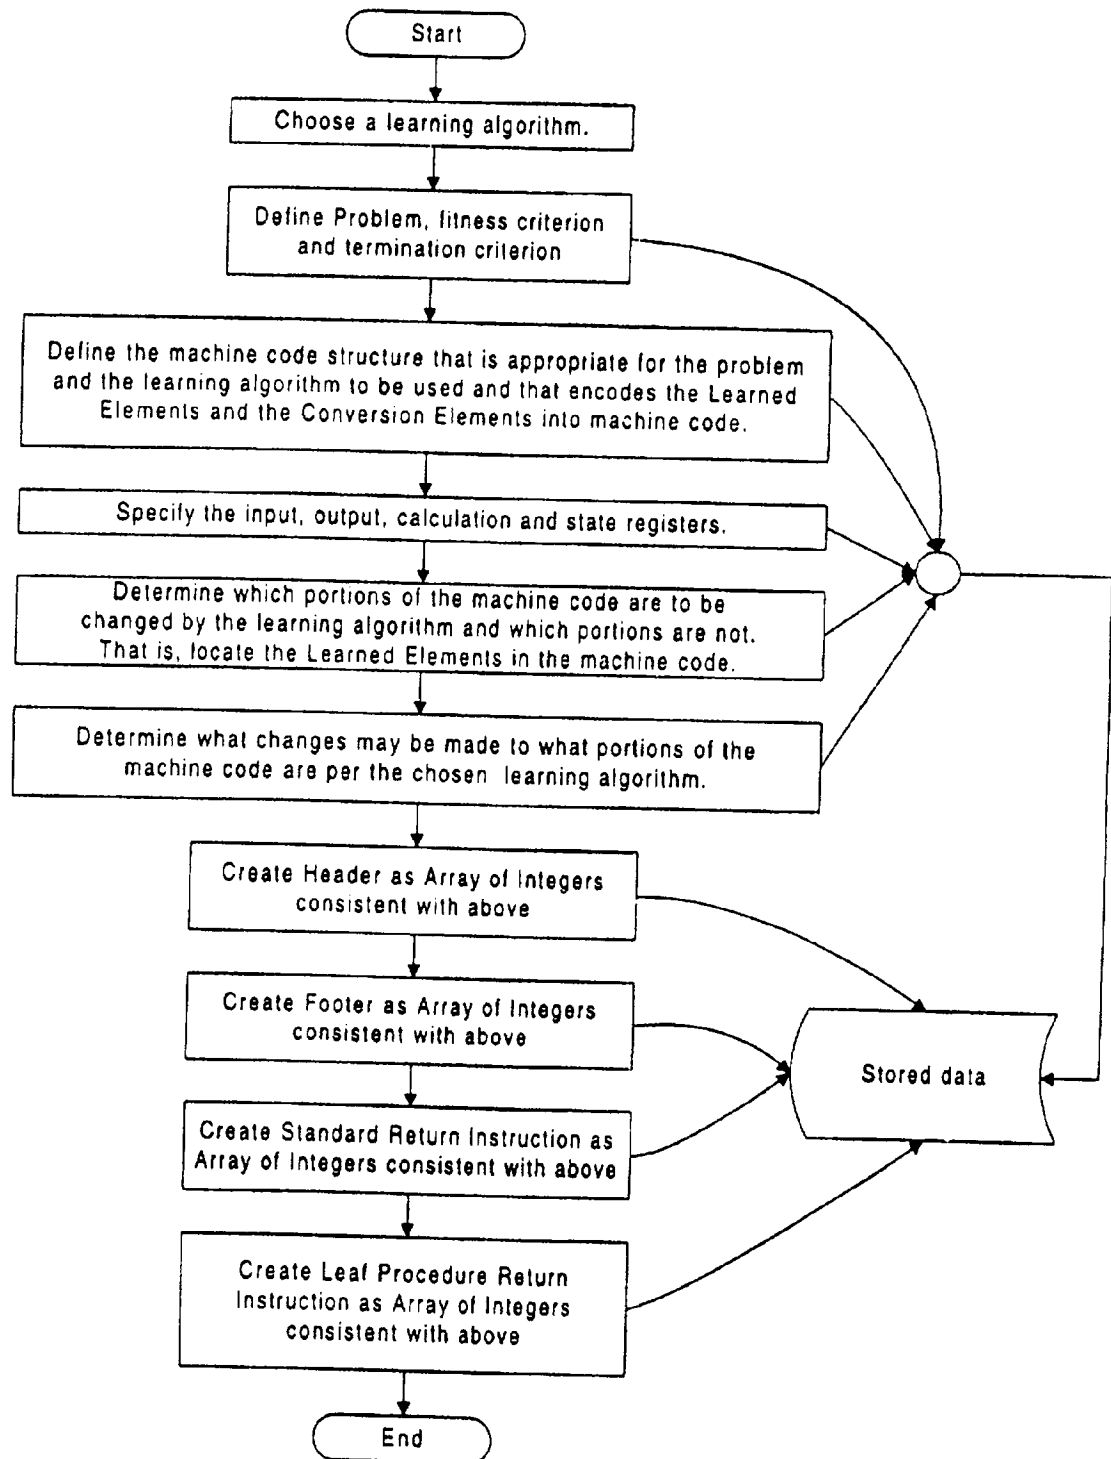


Figure 16b

Details of
system Definition
15/46

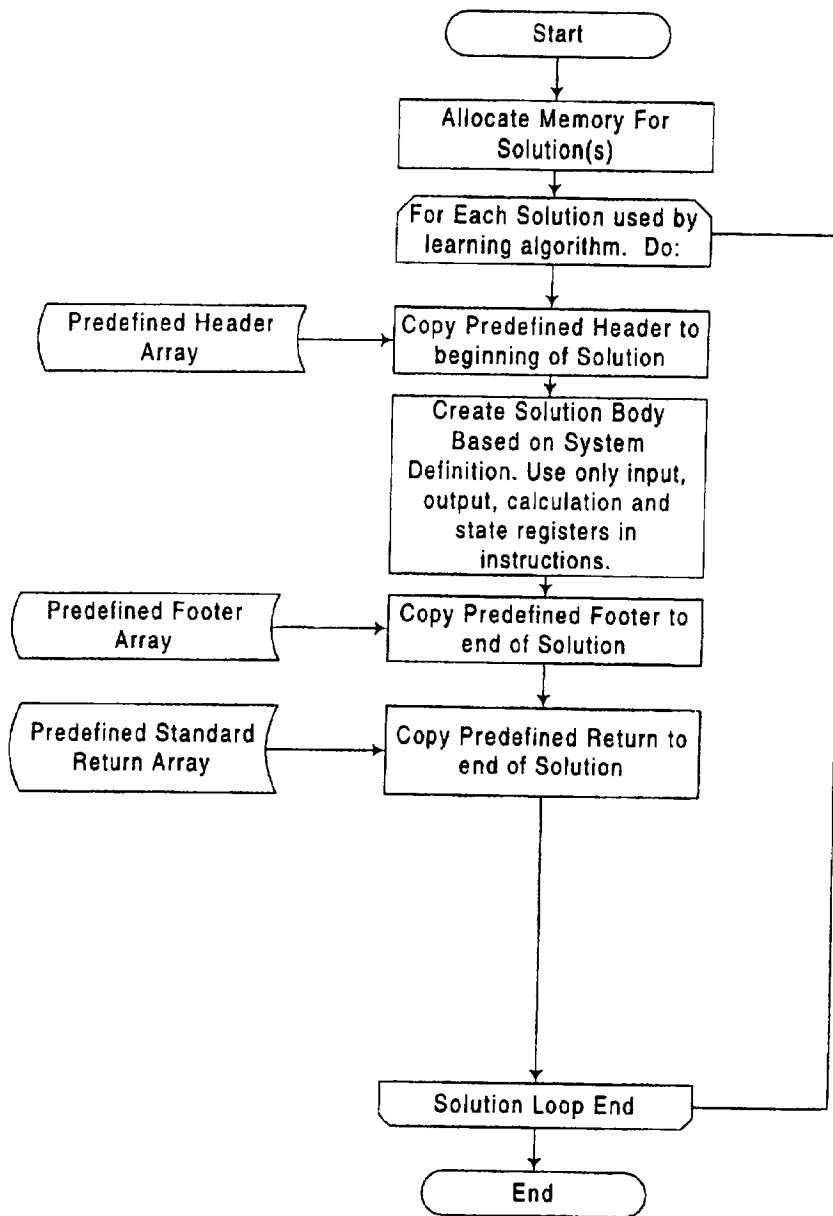
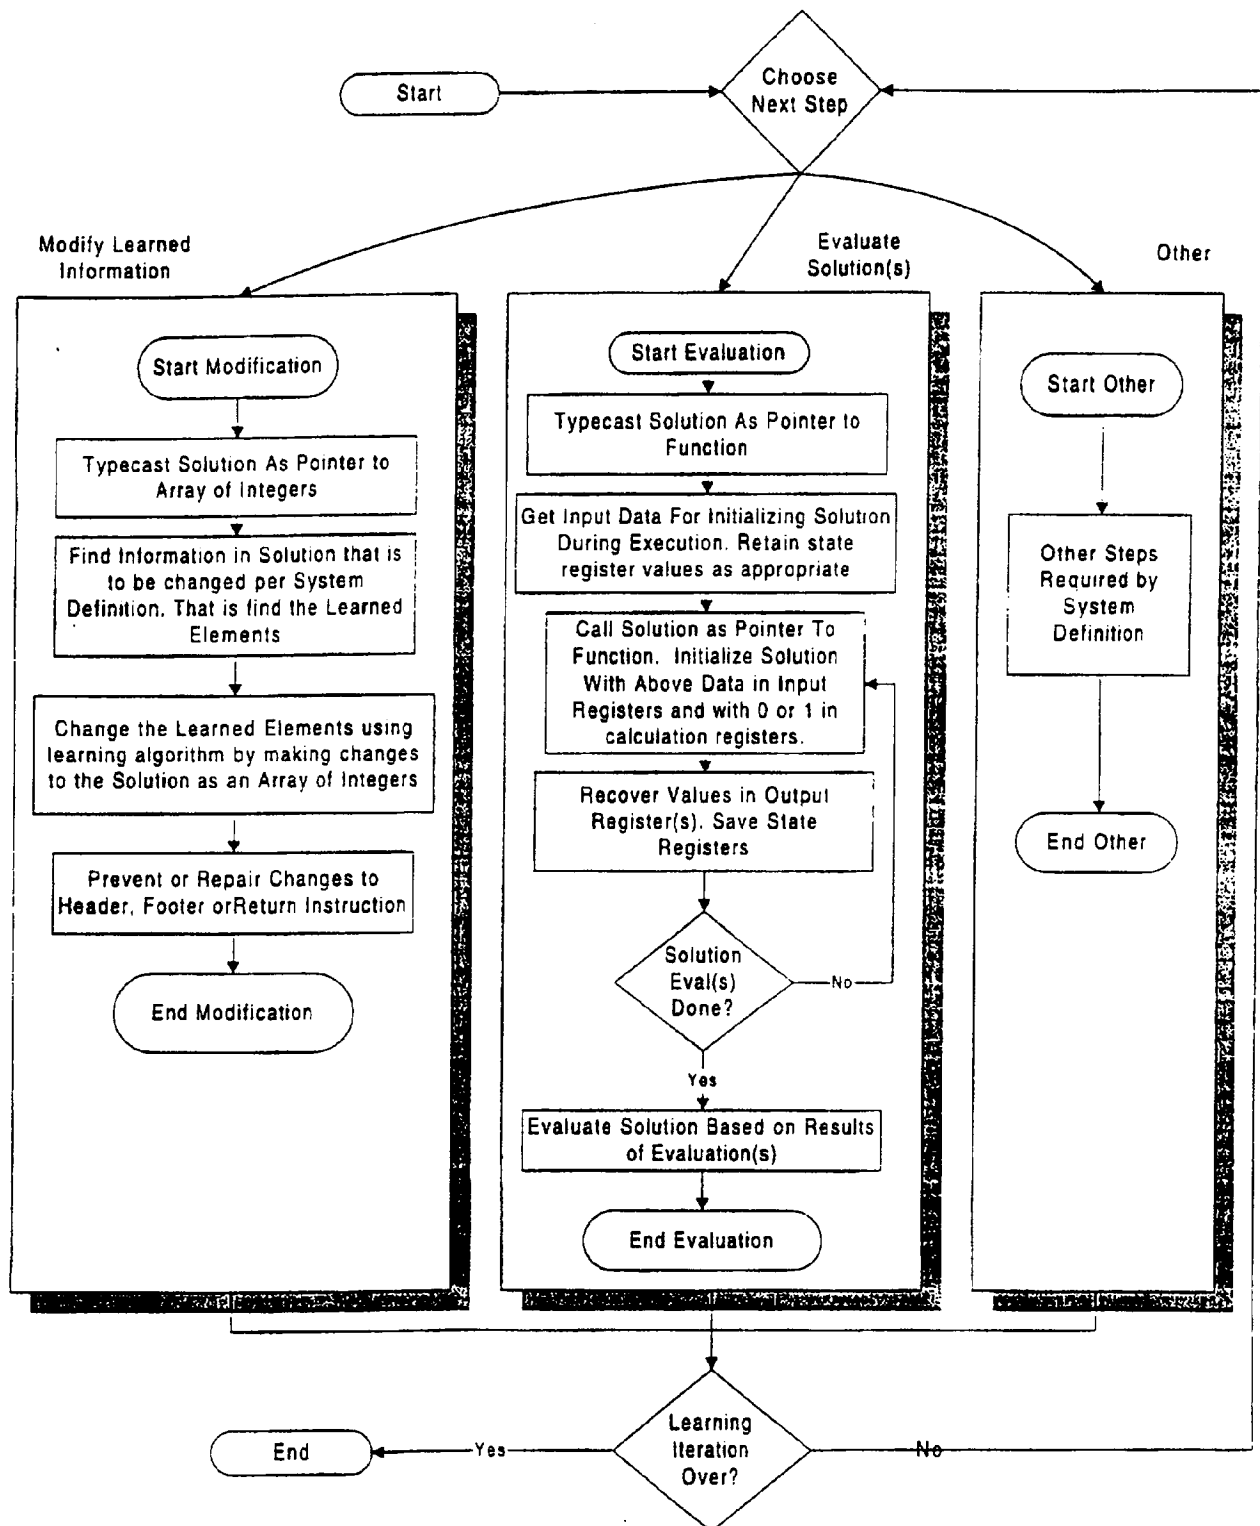


Figure 16c

Initialization

"Solution" means entity or individual

16/46



Solution means entity or individual

Figure 16d

Learn for One Iteration

17/46

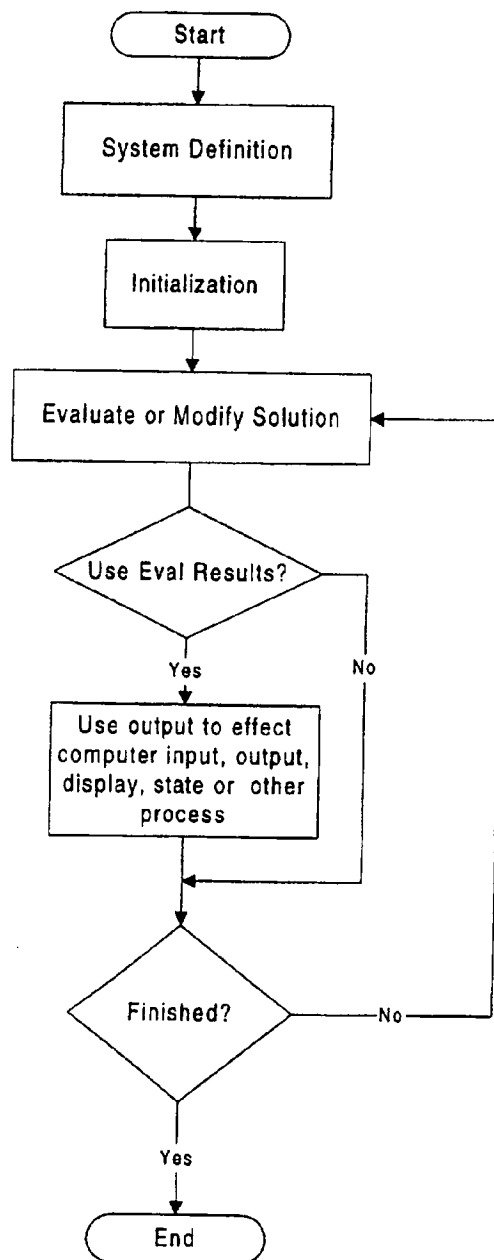


Figure 17a

General Operation of the Invention As Repetitive Computation System

"Solution" means entity or individual

18/46

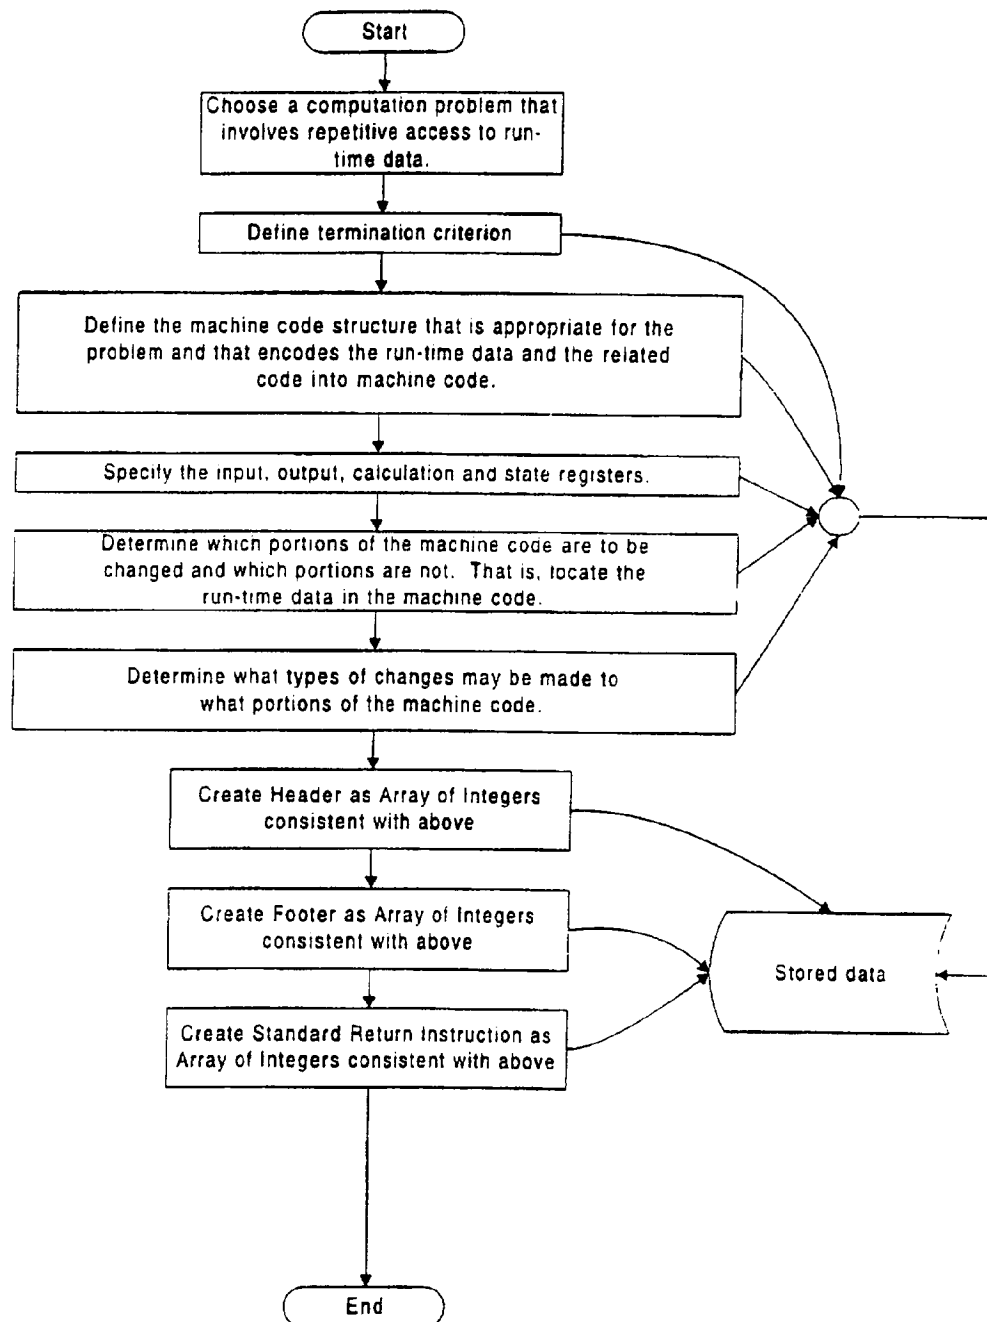


Figure 17b

System Definition

Solution means entity or individual

19/46

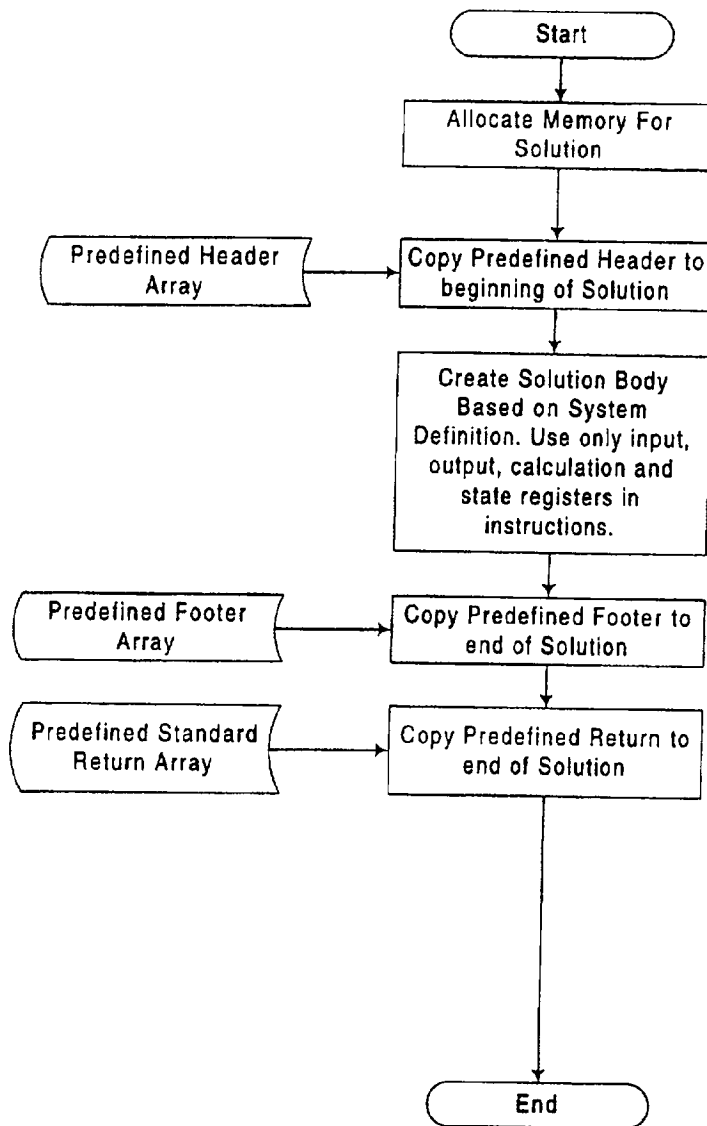
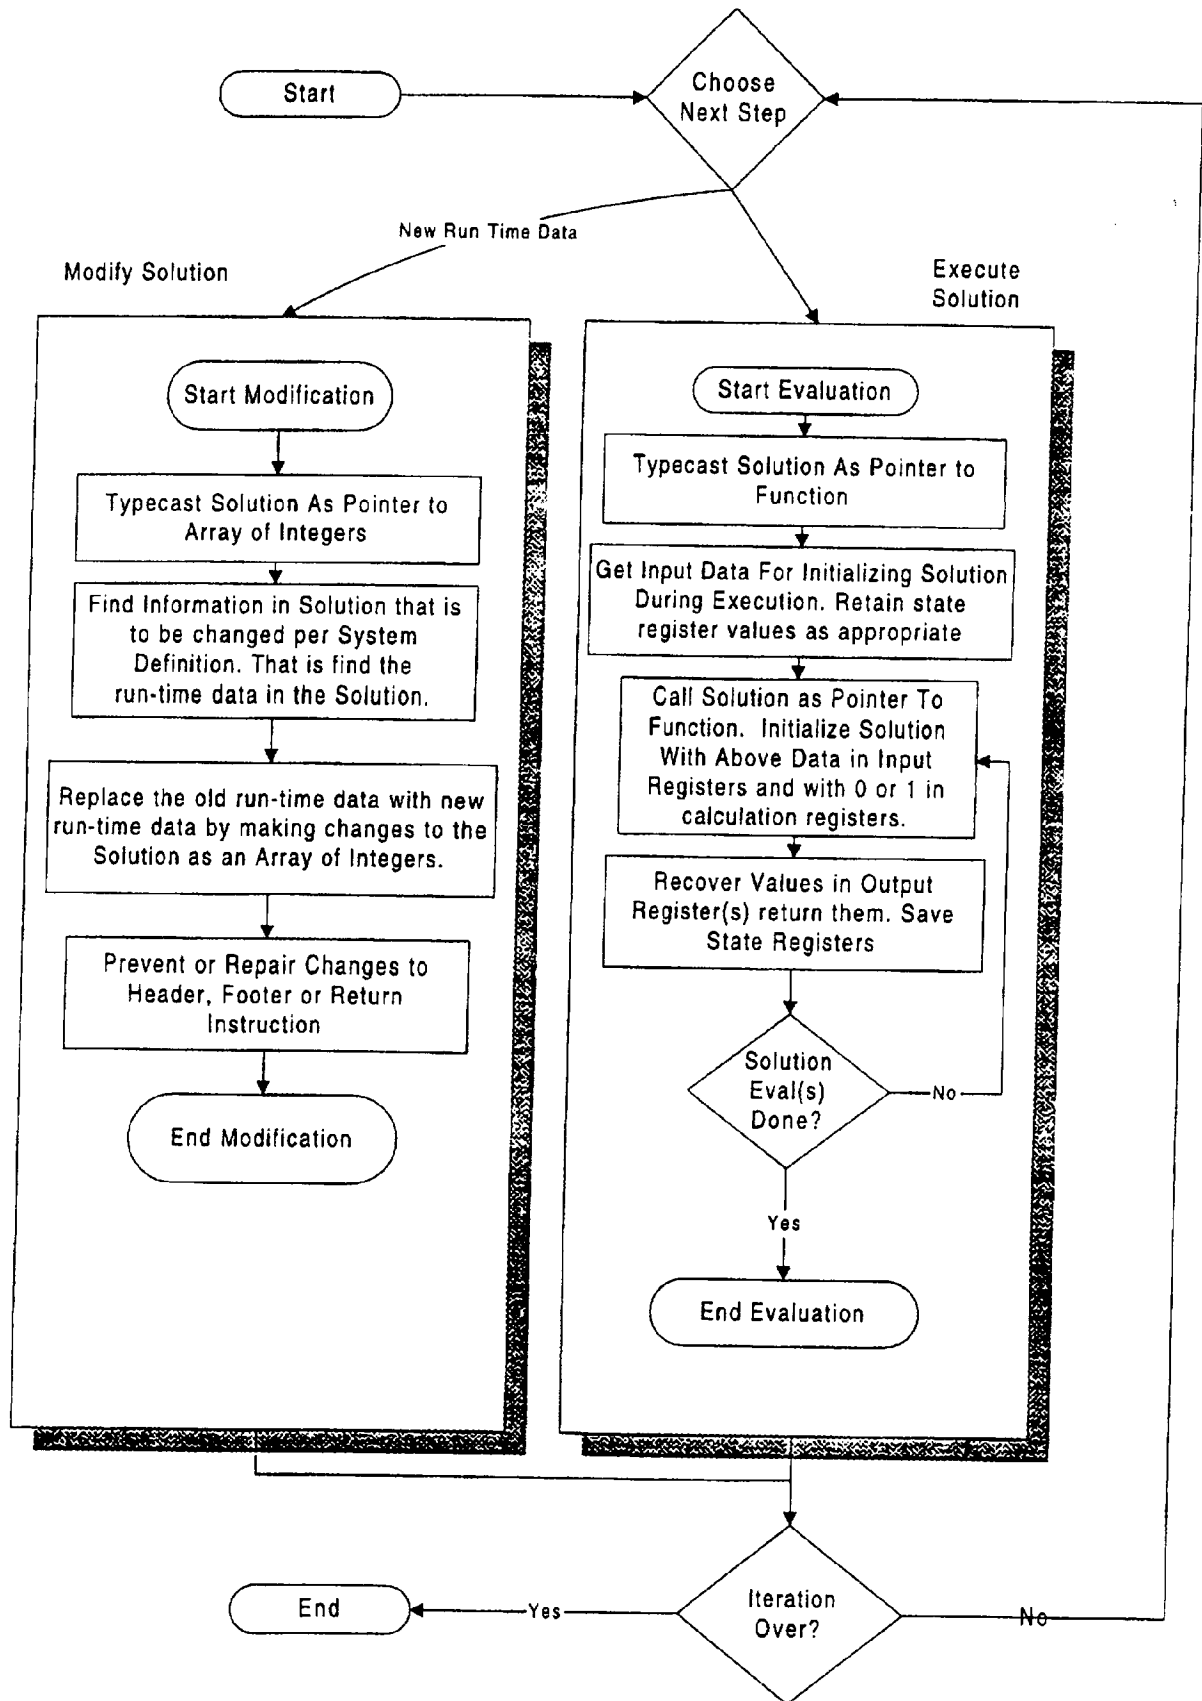


Figure 17c

Initialization

"Solution" means entity or individual

20/46



Solution means entity or individual

Figure 17d

Evaluate or Modify Function

21/46

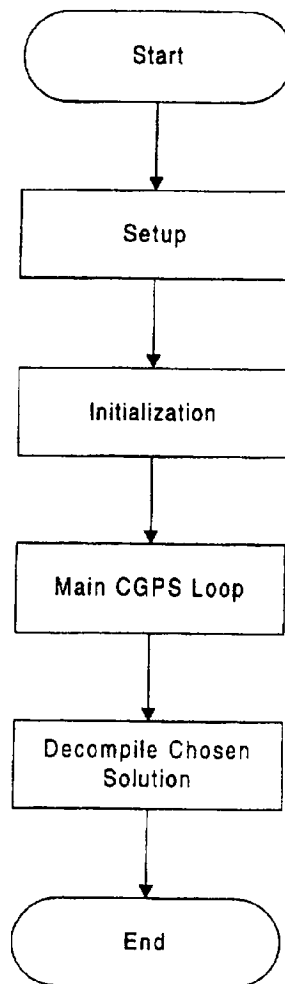


Figure 18a

22/46

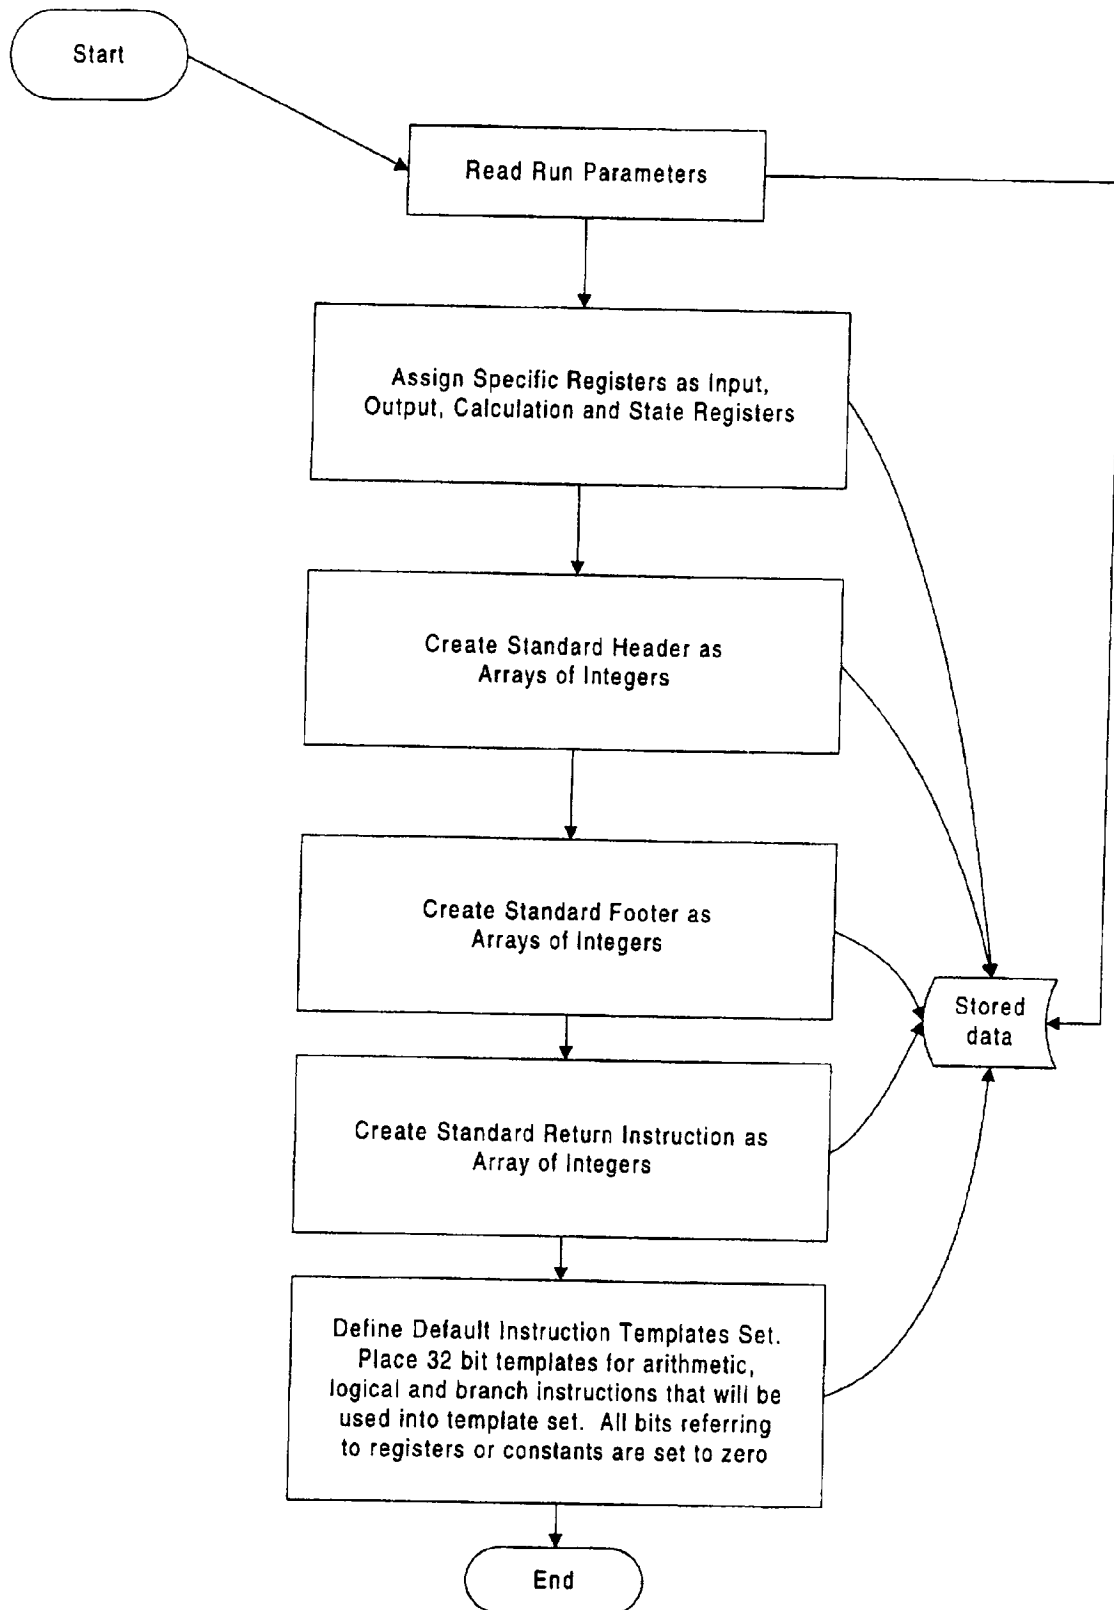


Figure 18b

Setup
23/46

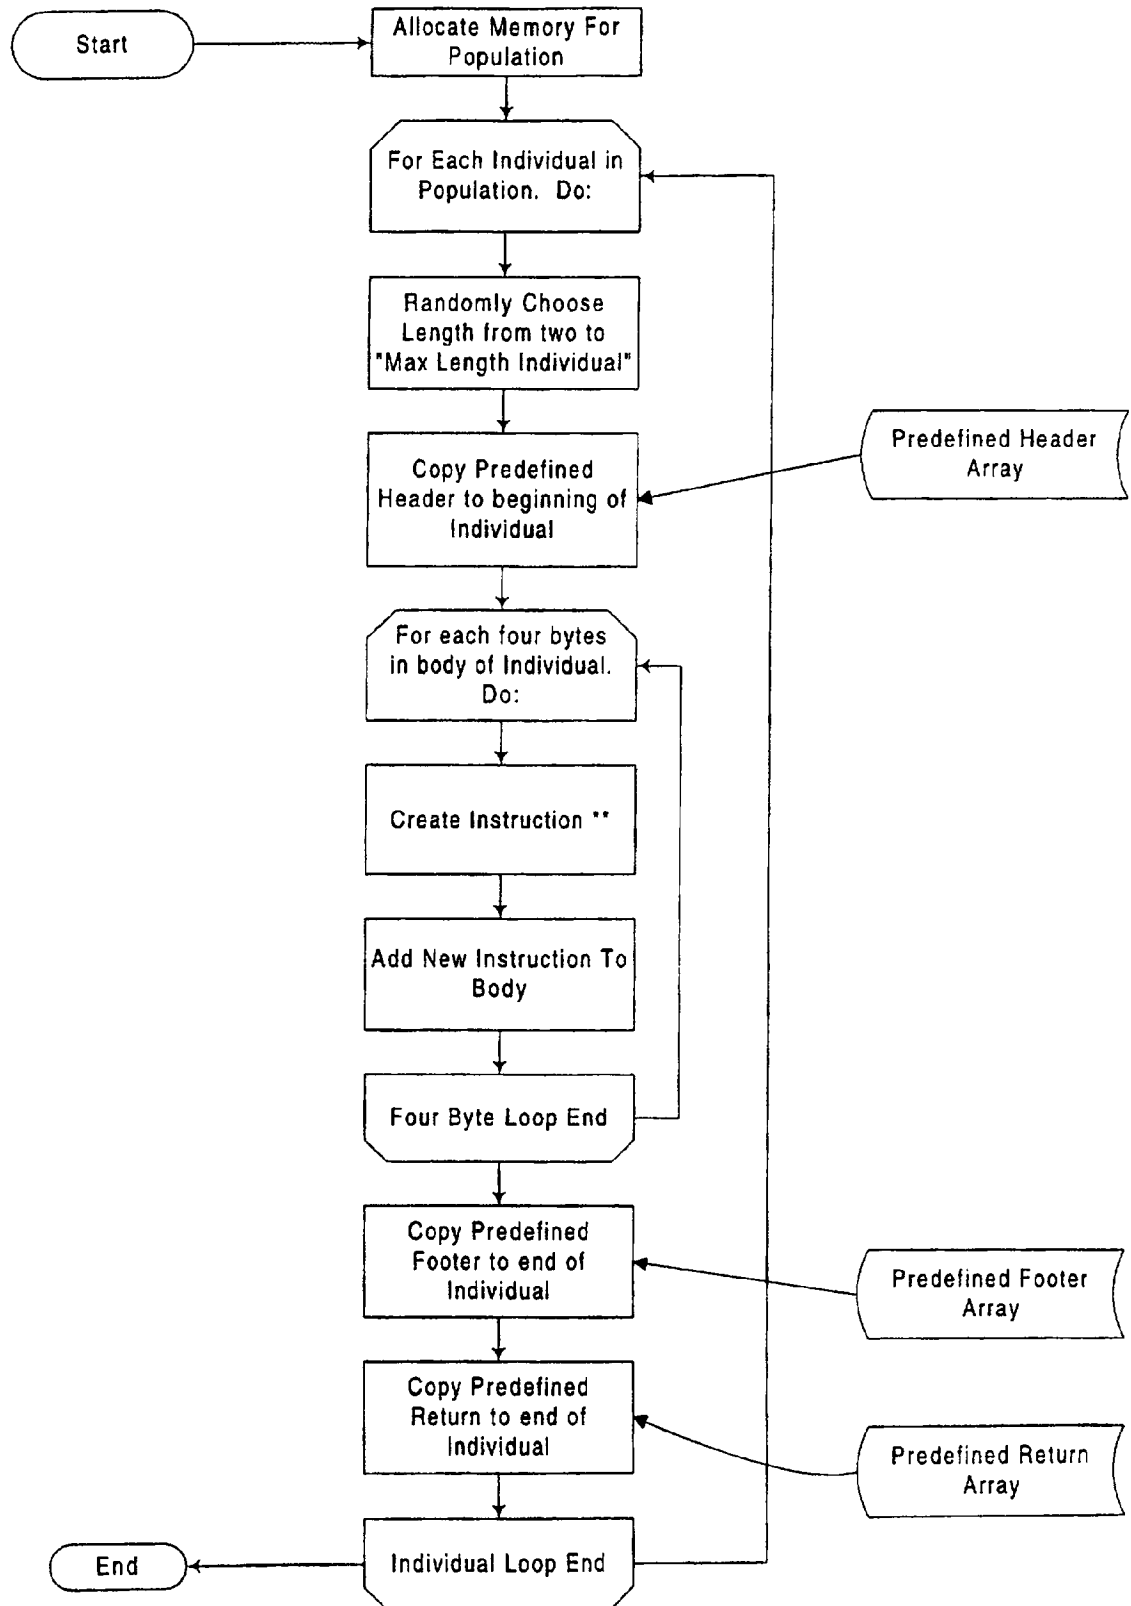


Figure 18c

Initialization
24/46

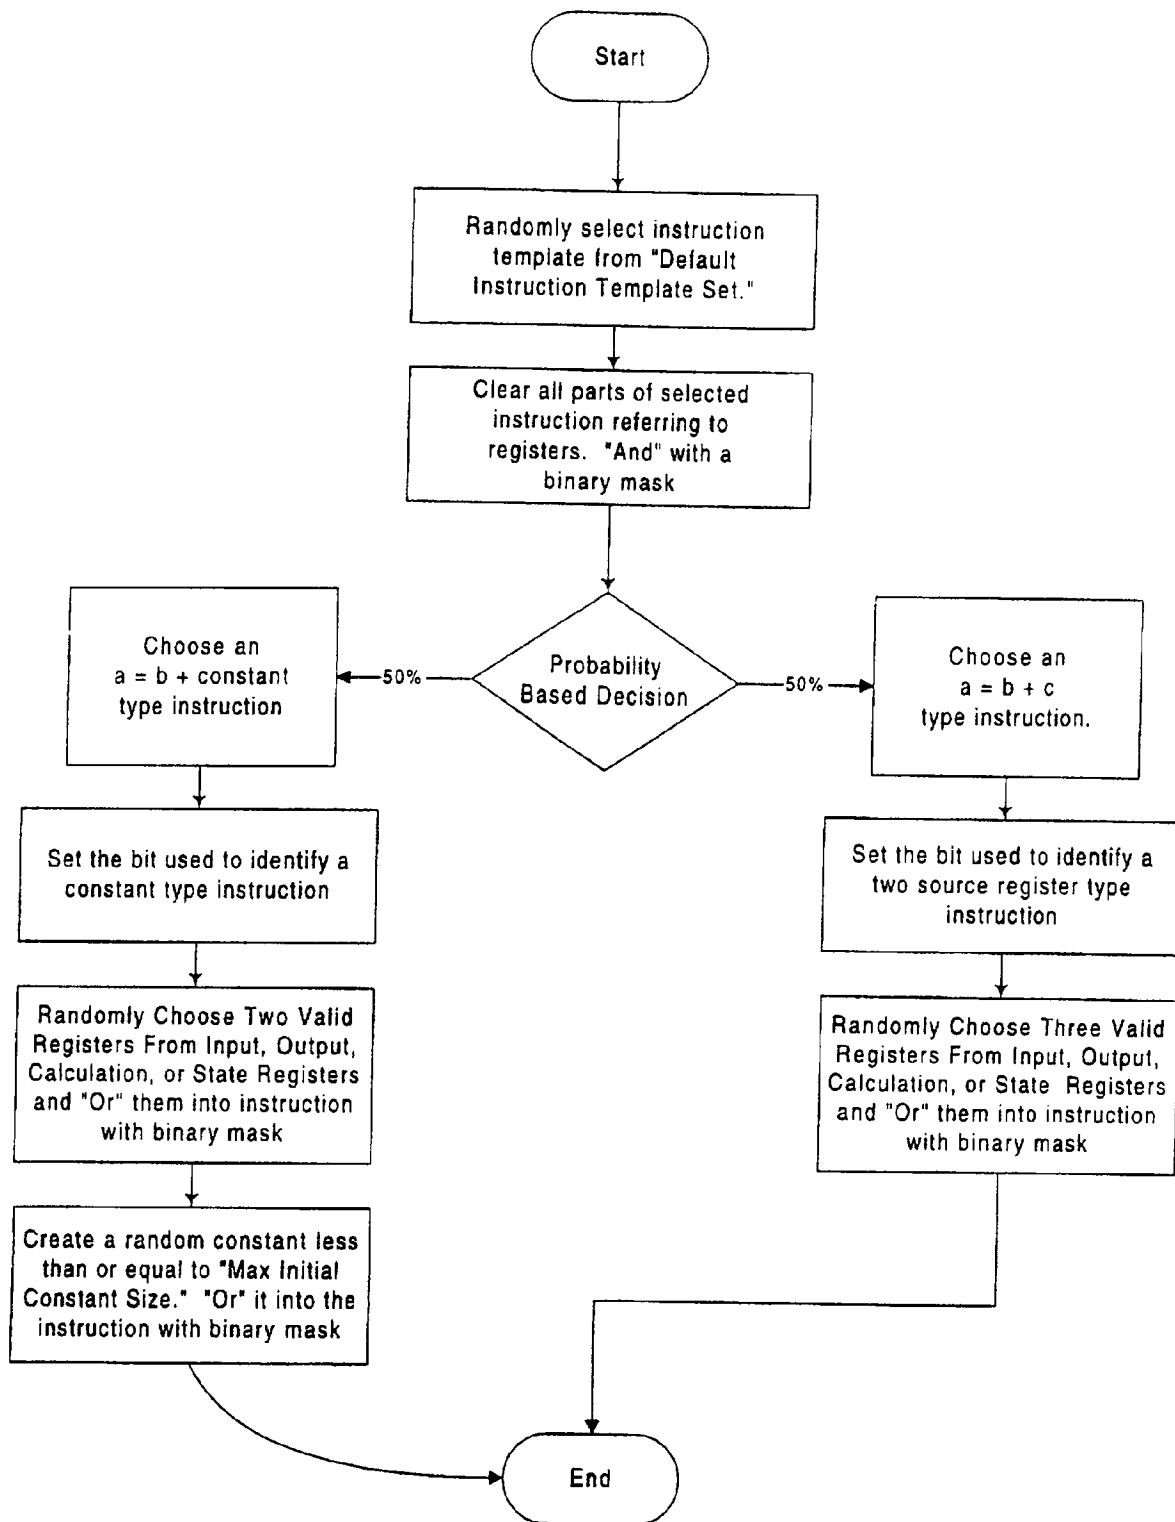


Figure 18d

Create Instruction

25/46

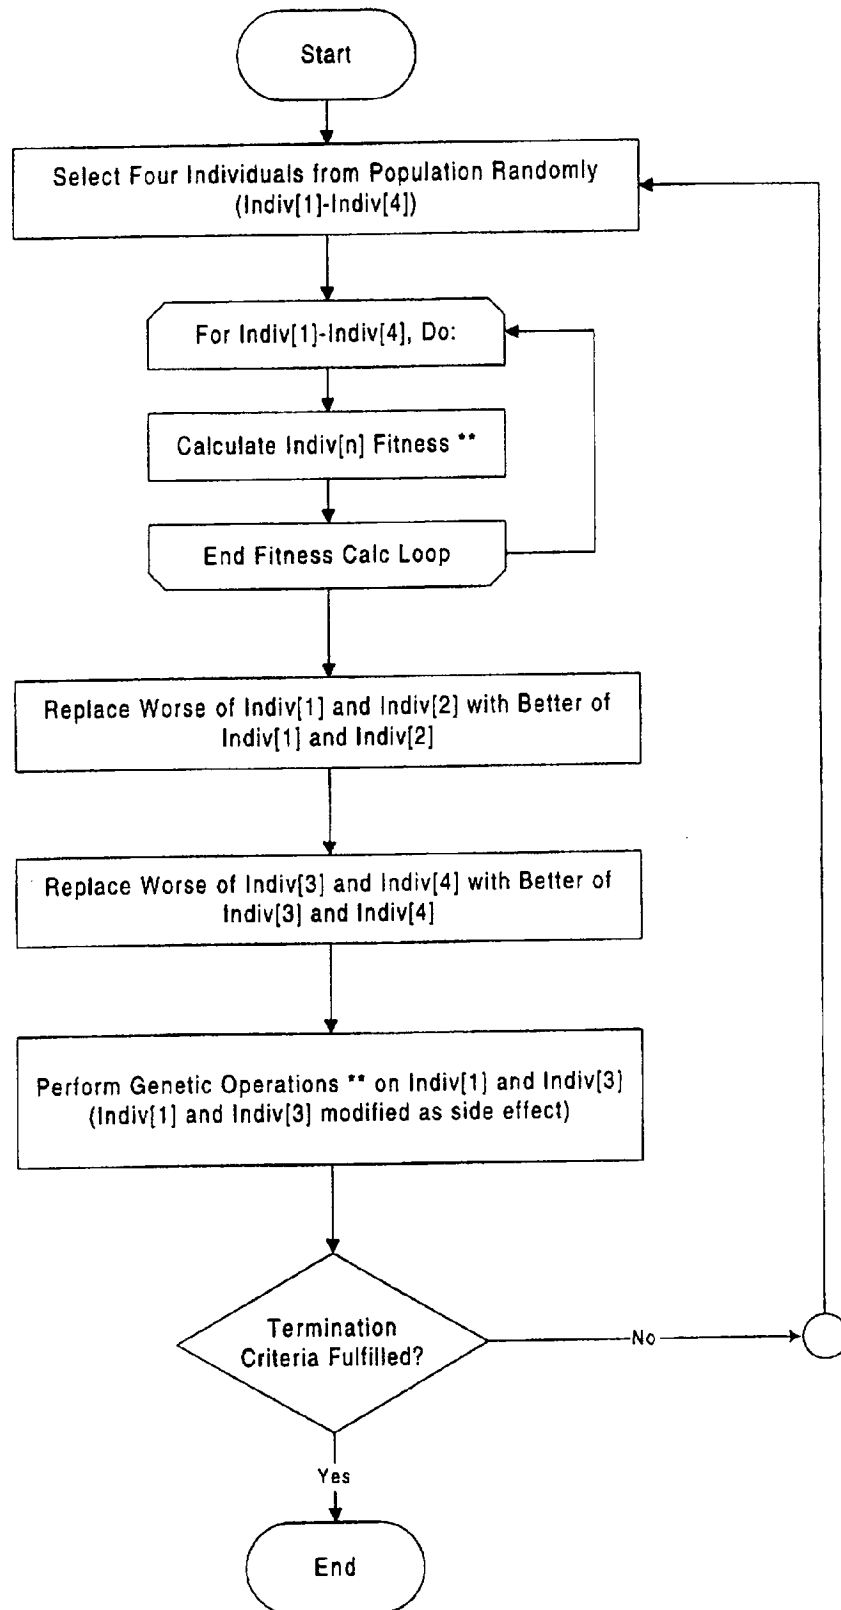
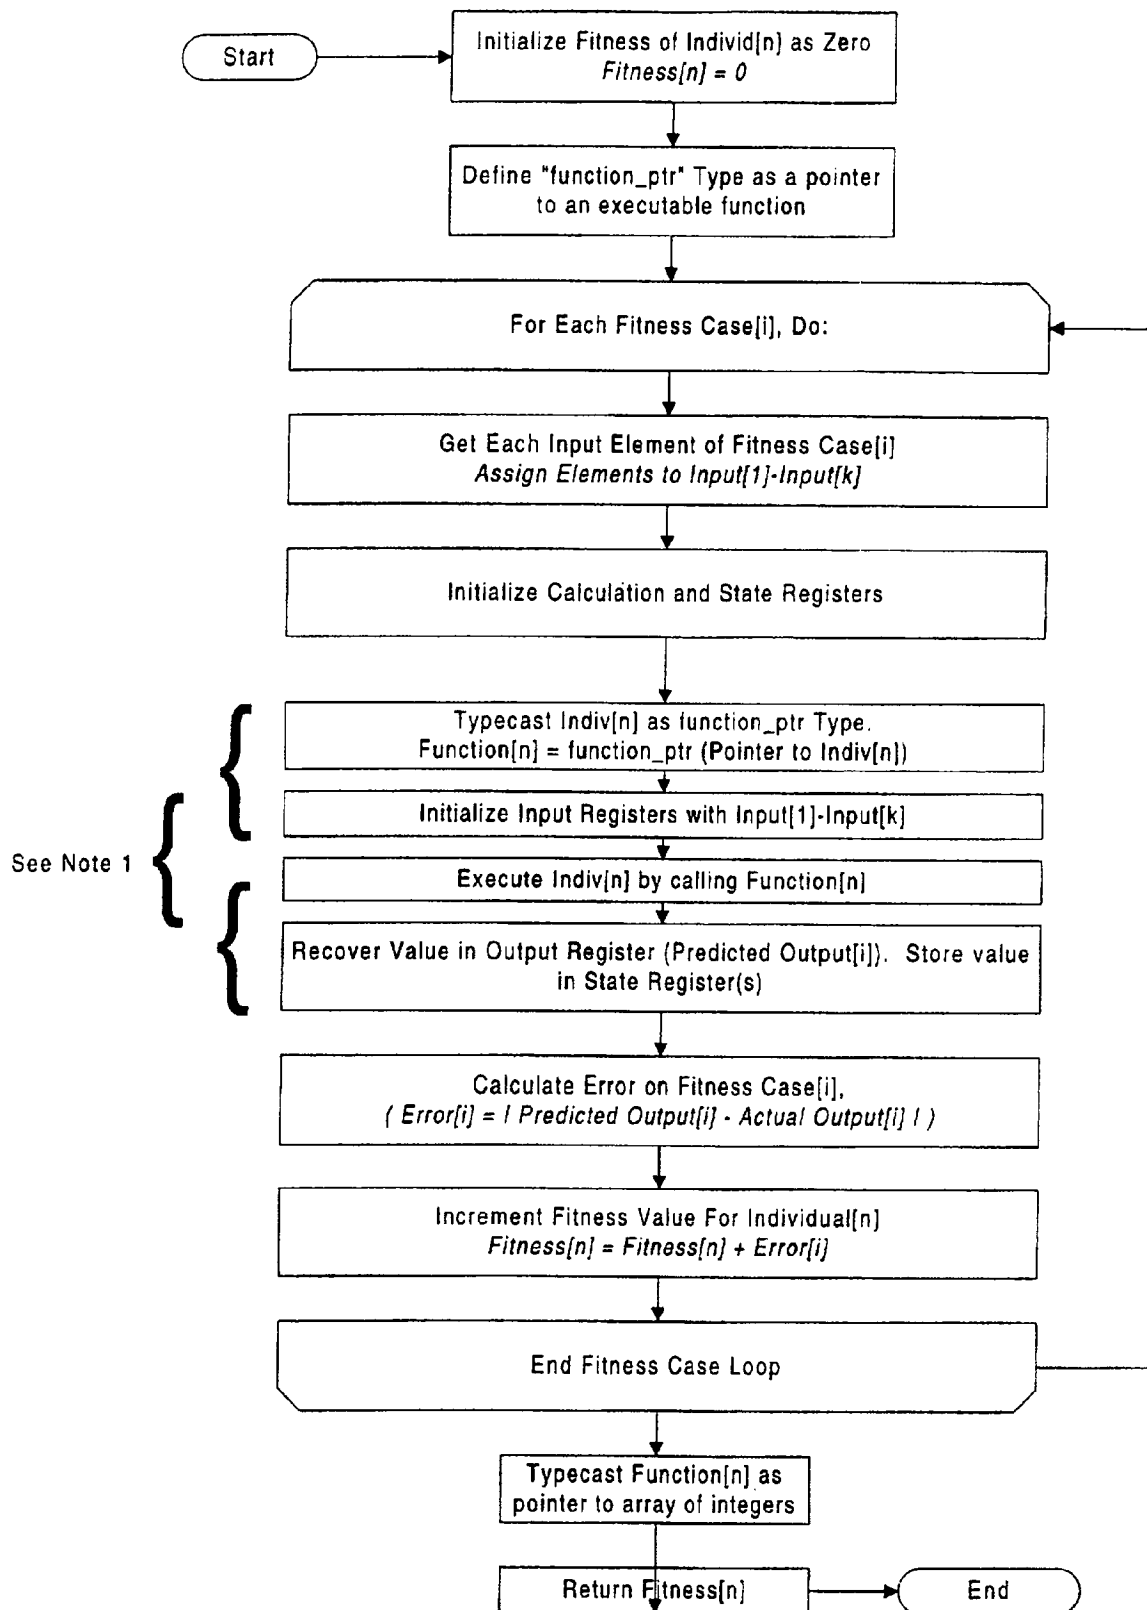


Figure 18e

Main CGPS Loop
26/46



Note 1: All of these steps are accomplished with the code:
 $\text{Predicted_Output}[i] = ((\text{function_ptr}) \text{Indiv}[n])(\text{Input}[1], \dots, \text{Input}[k])$

Figure 18f

Calculate Indiv[n] Fitness

27/46

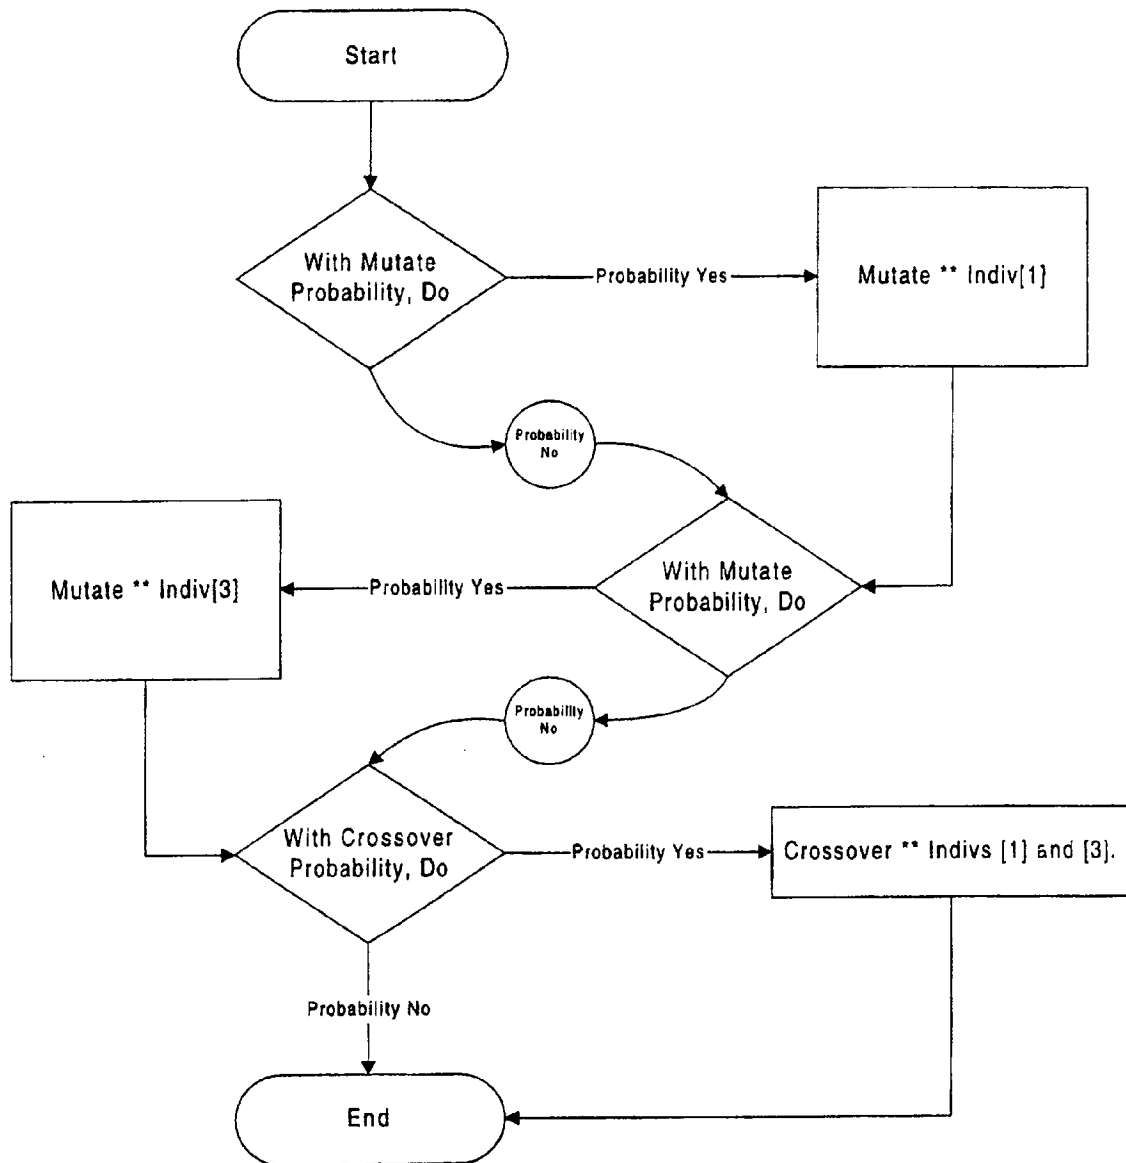


Figure 18g

Perform Genetic Operations

28/46

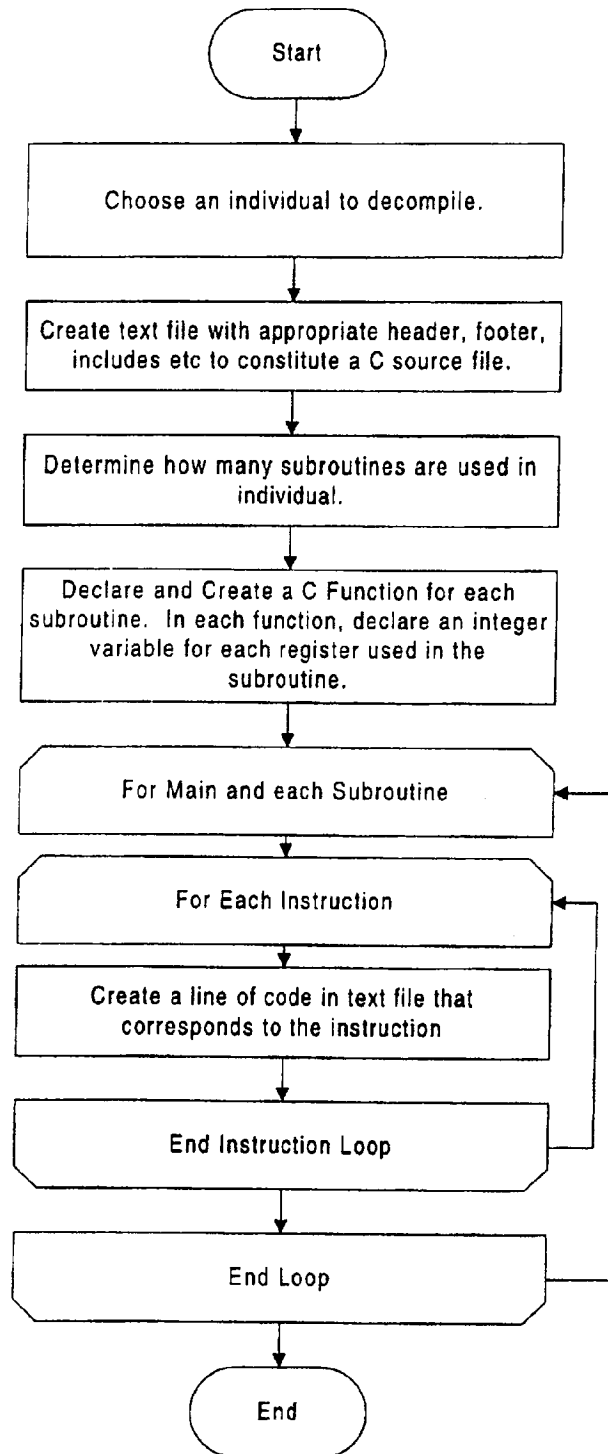


Figure 18h

29/46




Fig. 19

30/46

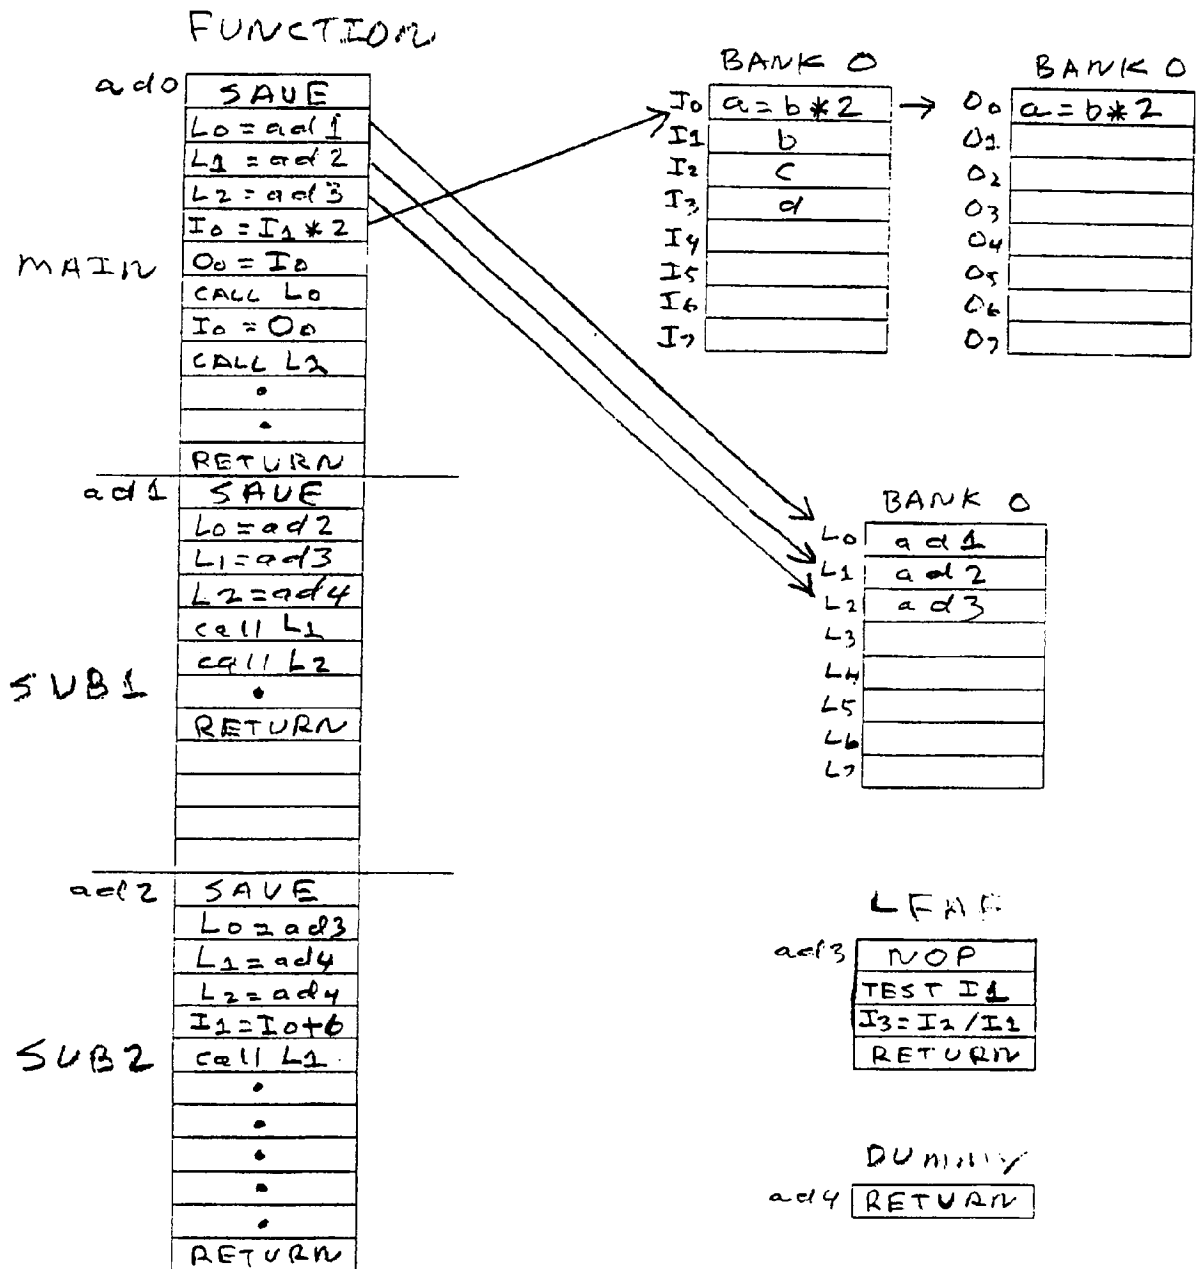


Fig. 20

31/46

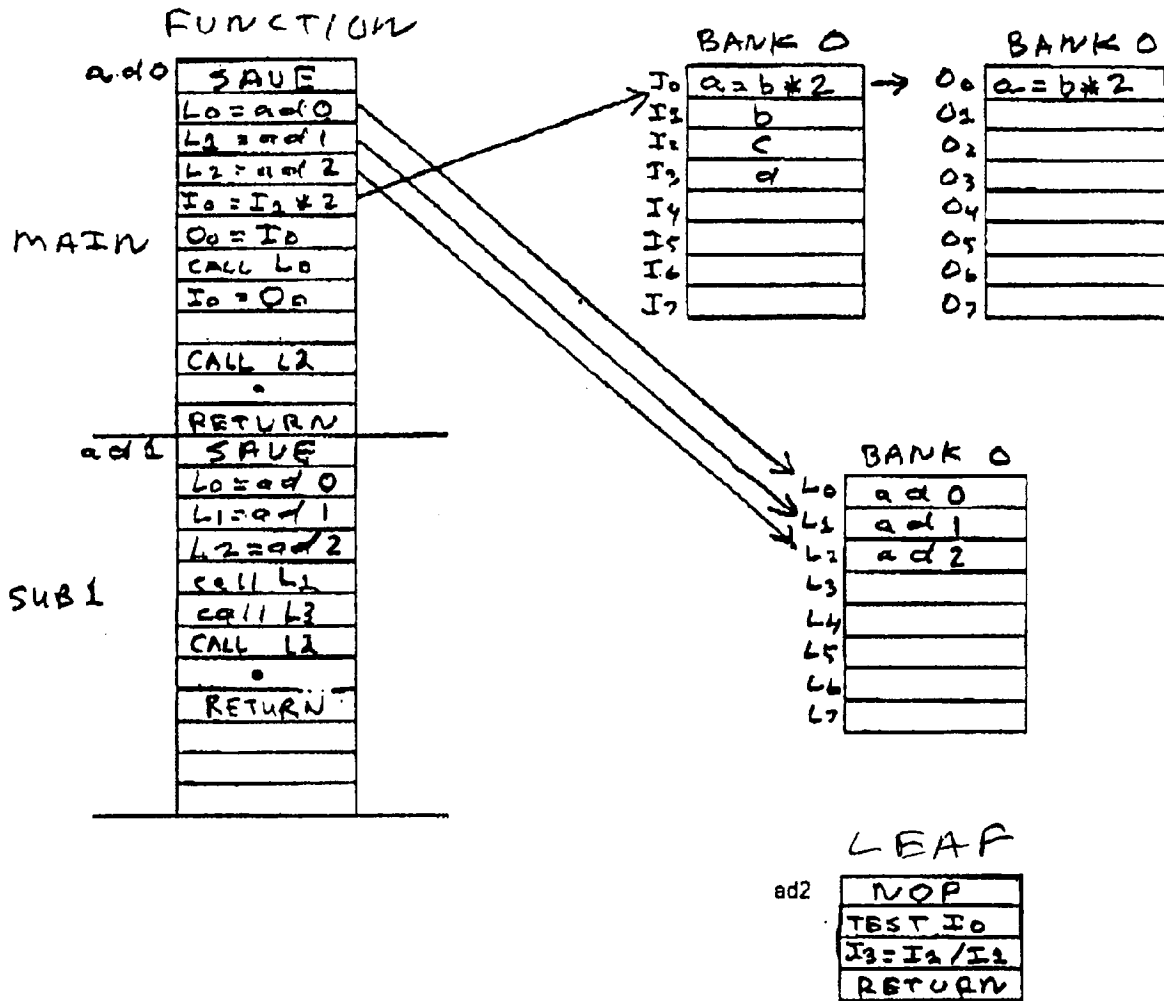


Fig. 21

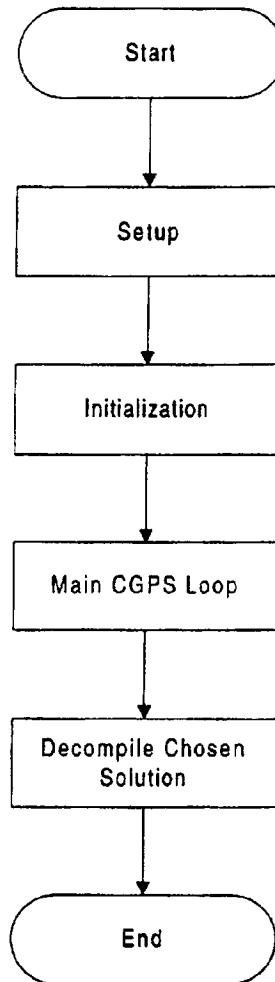


Figure 22a

33/46

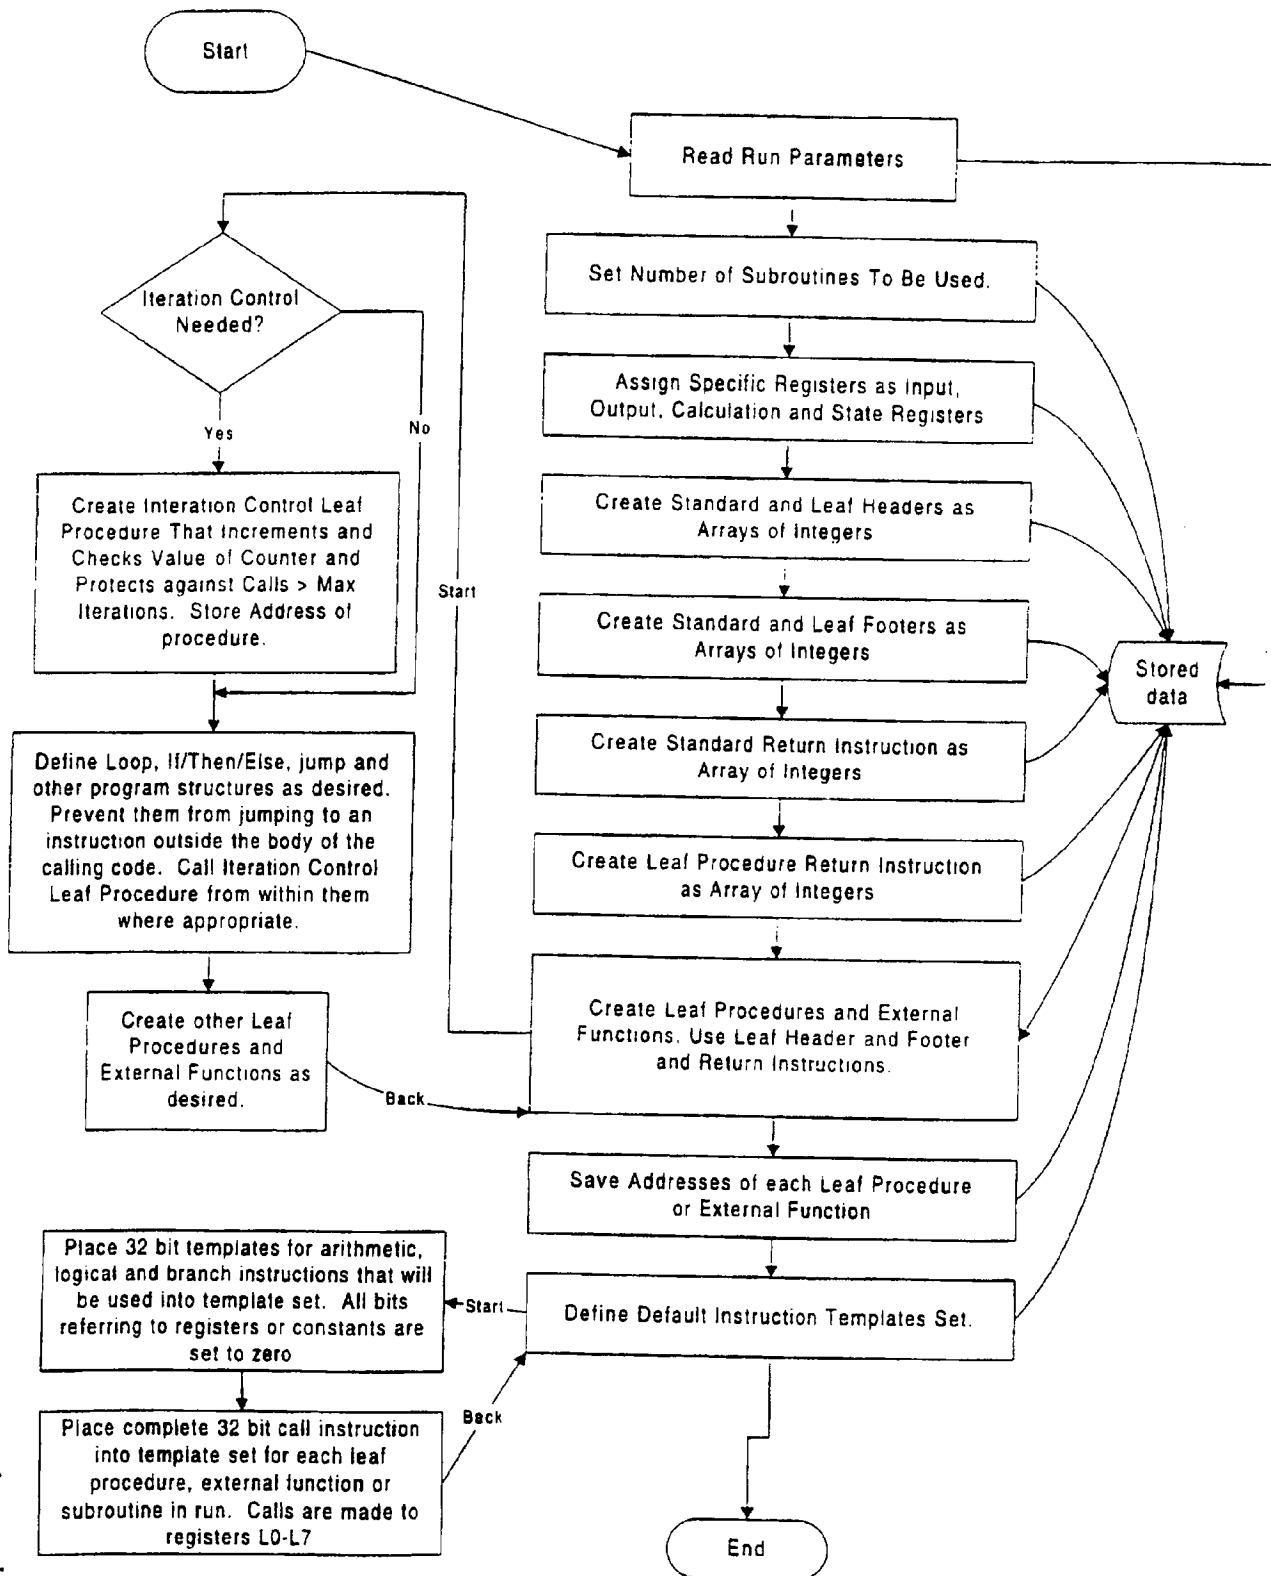


Figure 22b

34/46

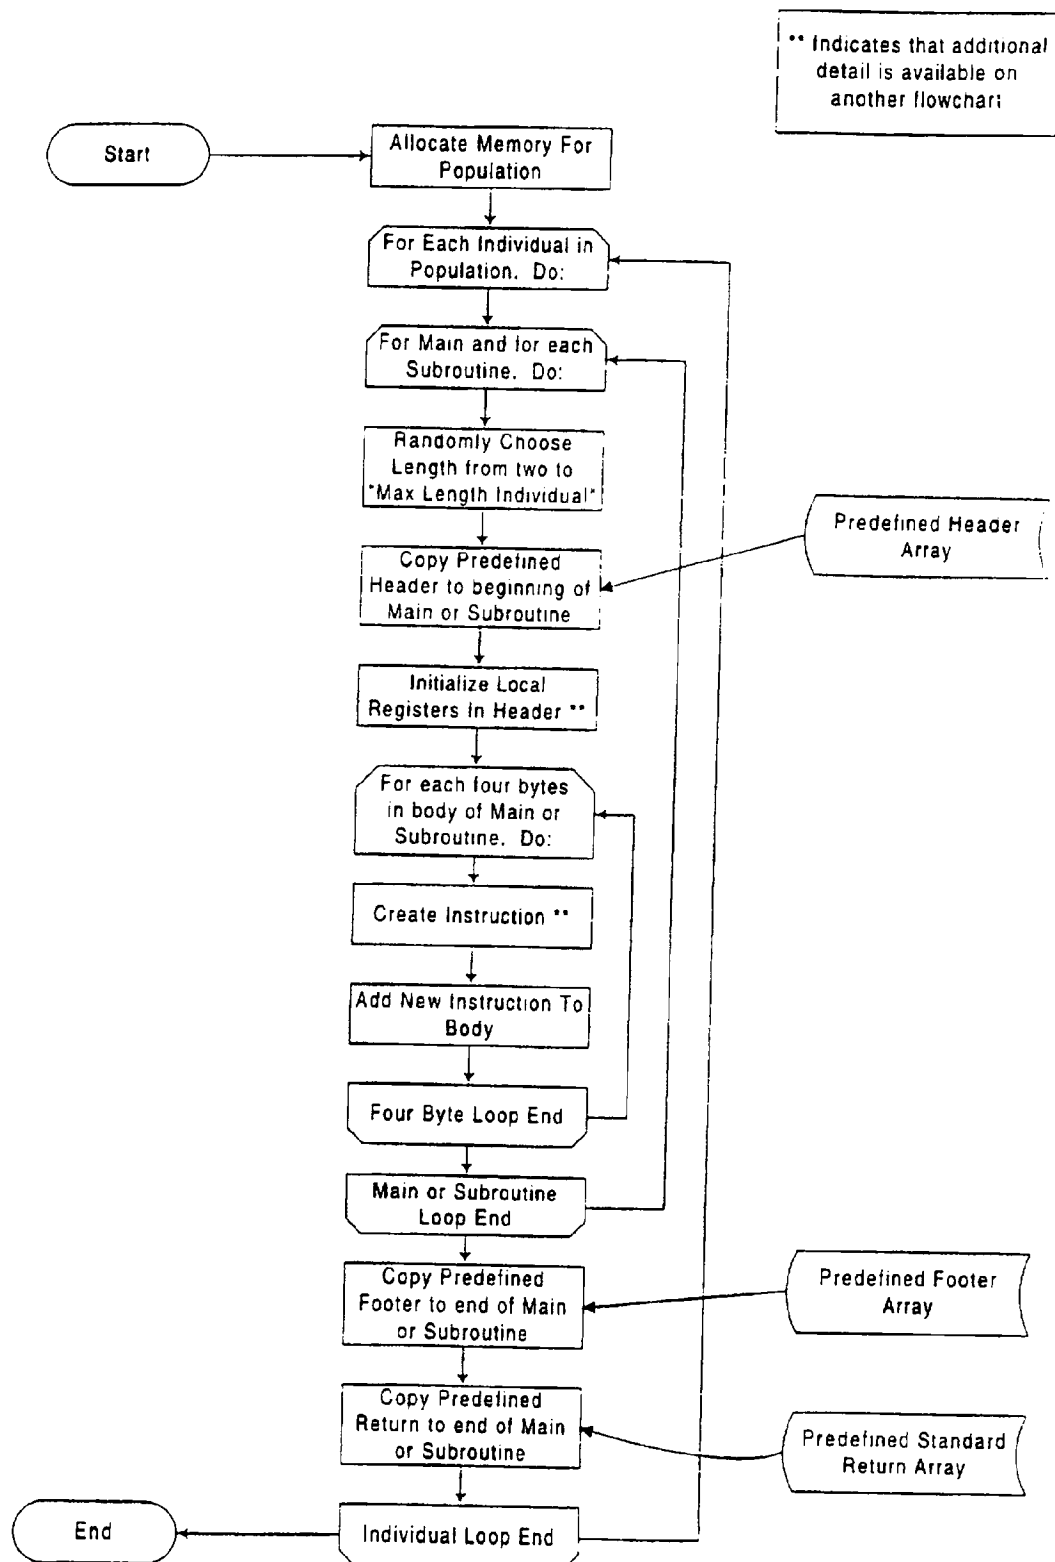


Figure 22c

35/46

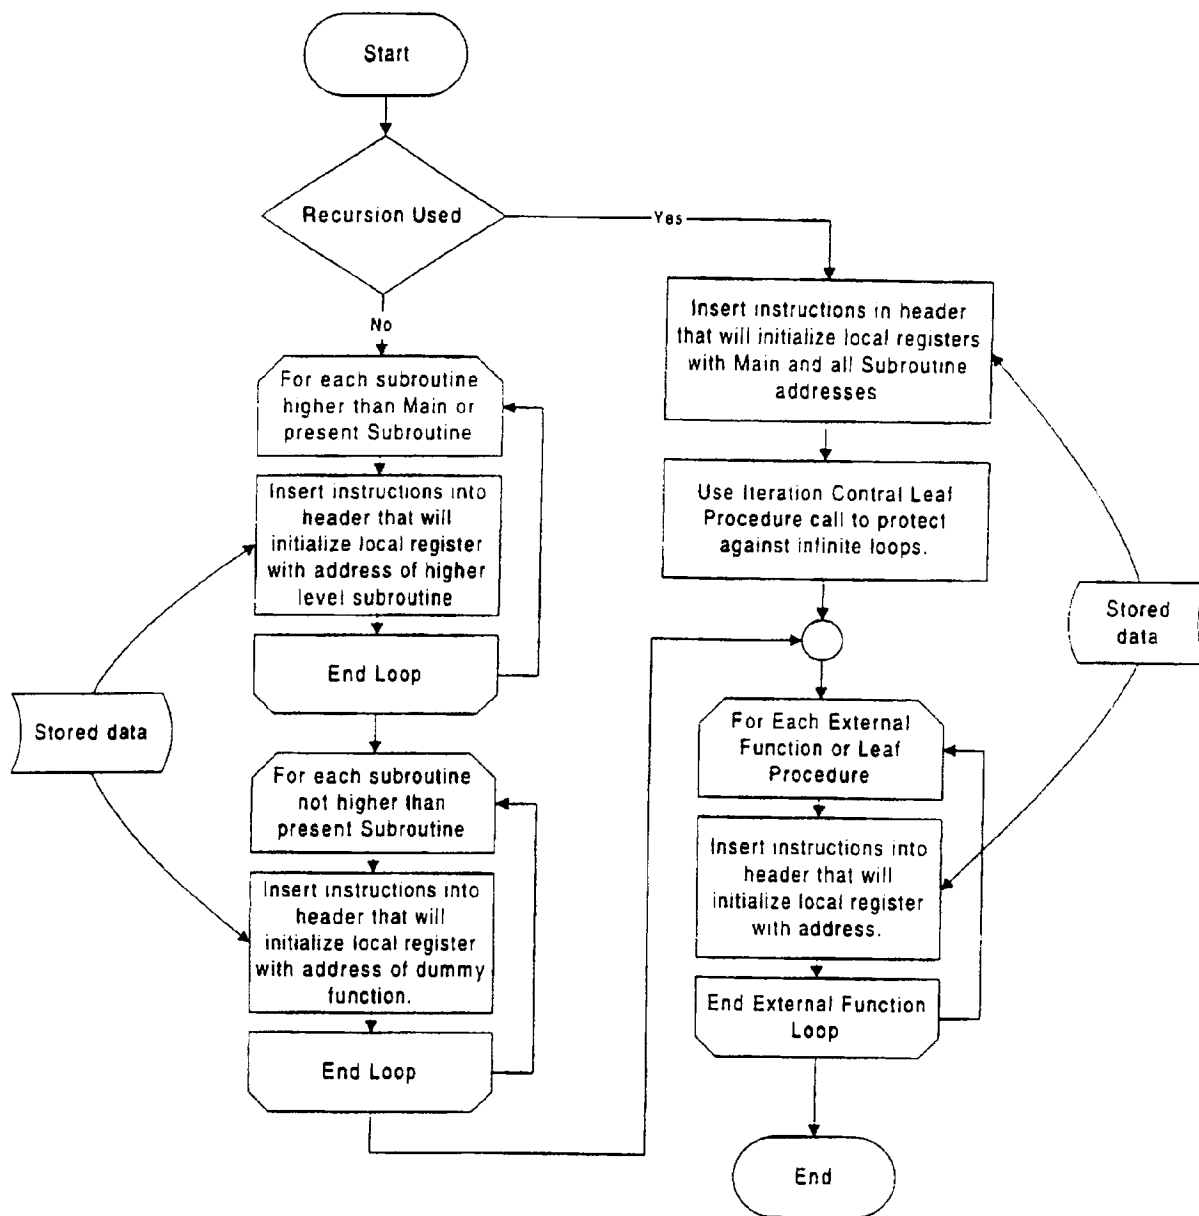


Figure 22d

36/46

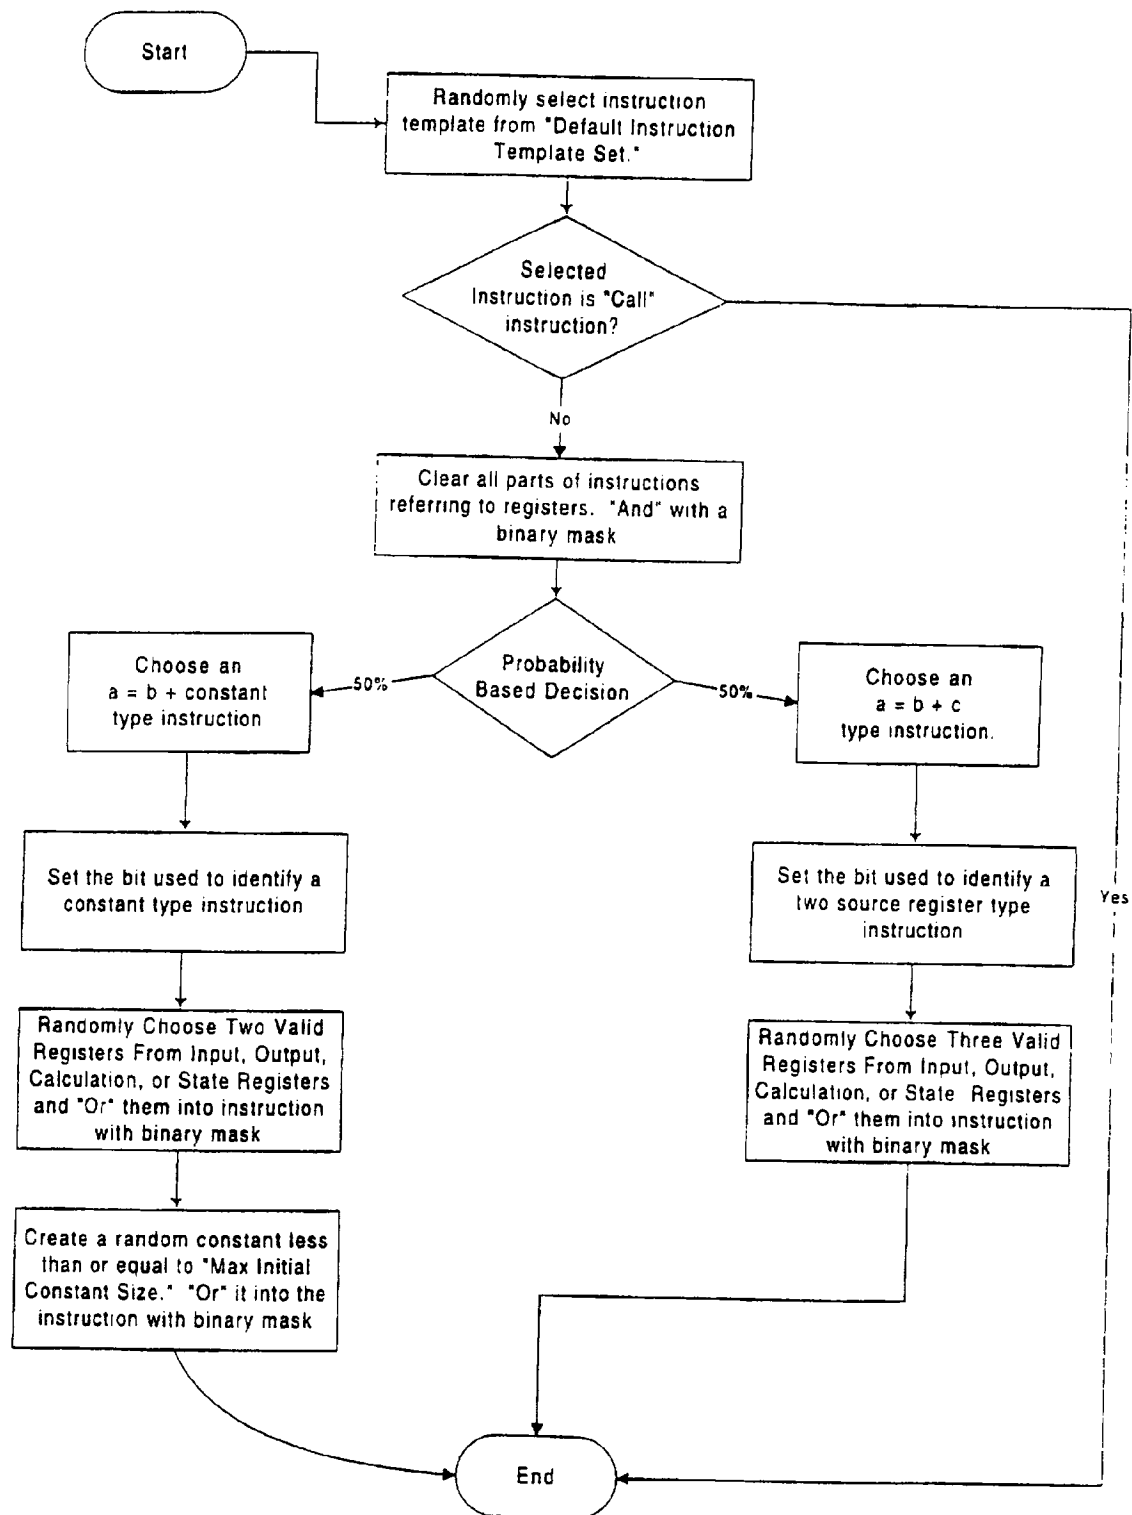


Figure 22e

37/46

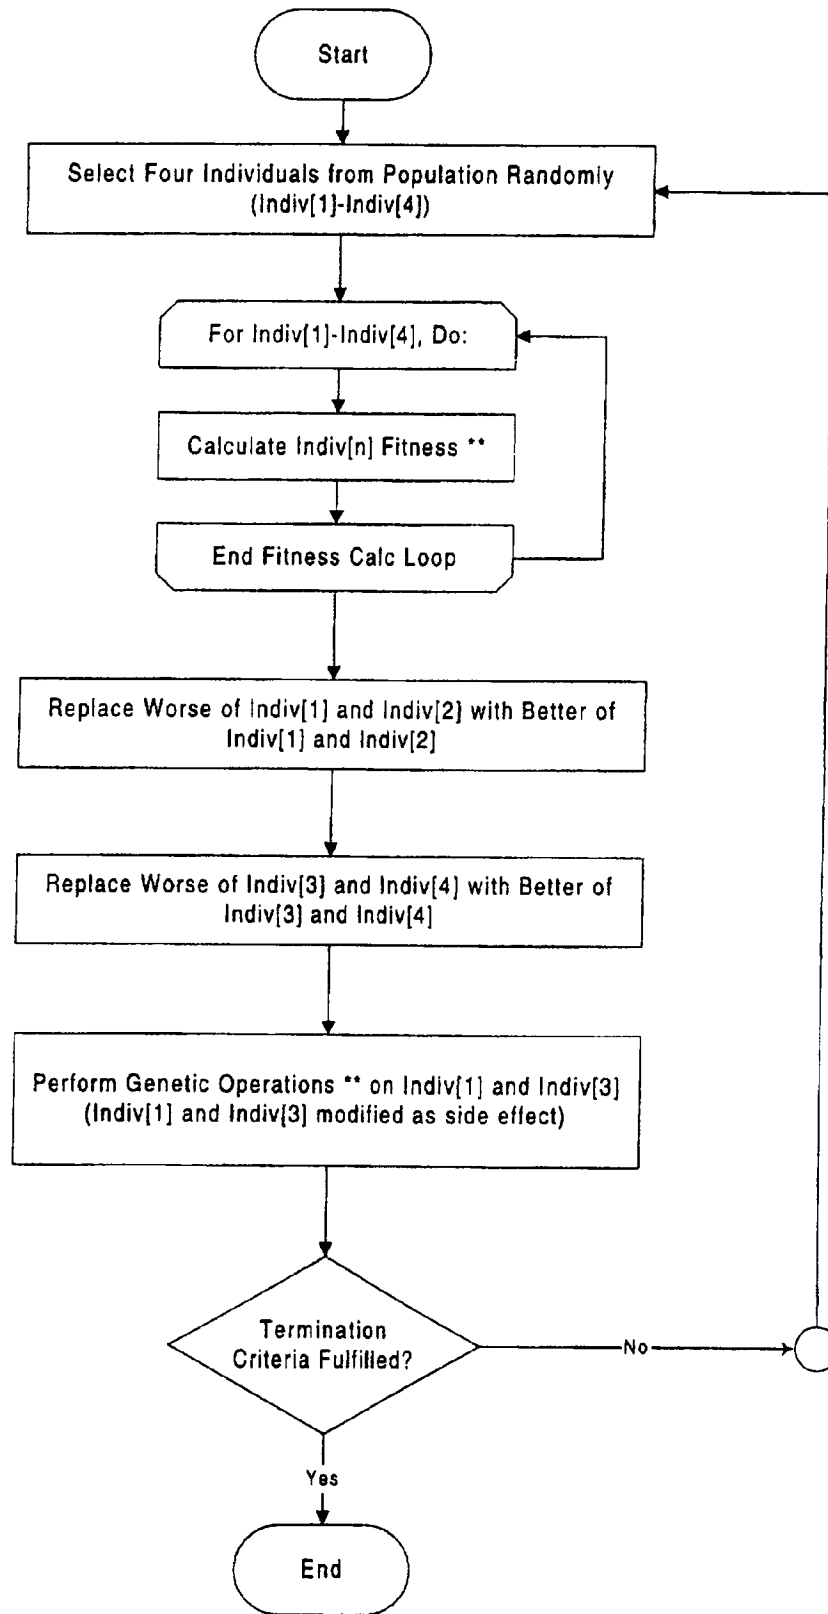


Figure 22f

38/46

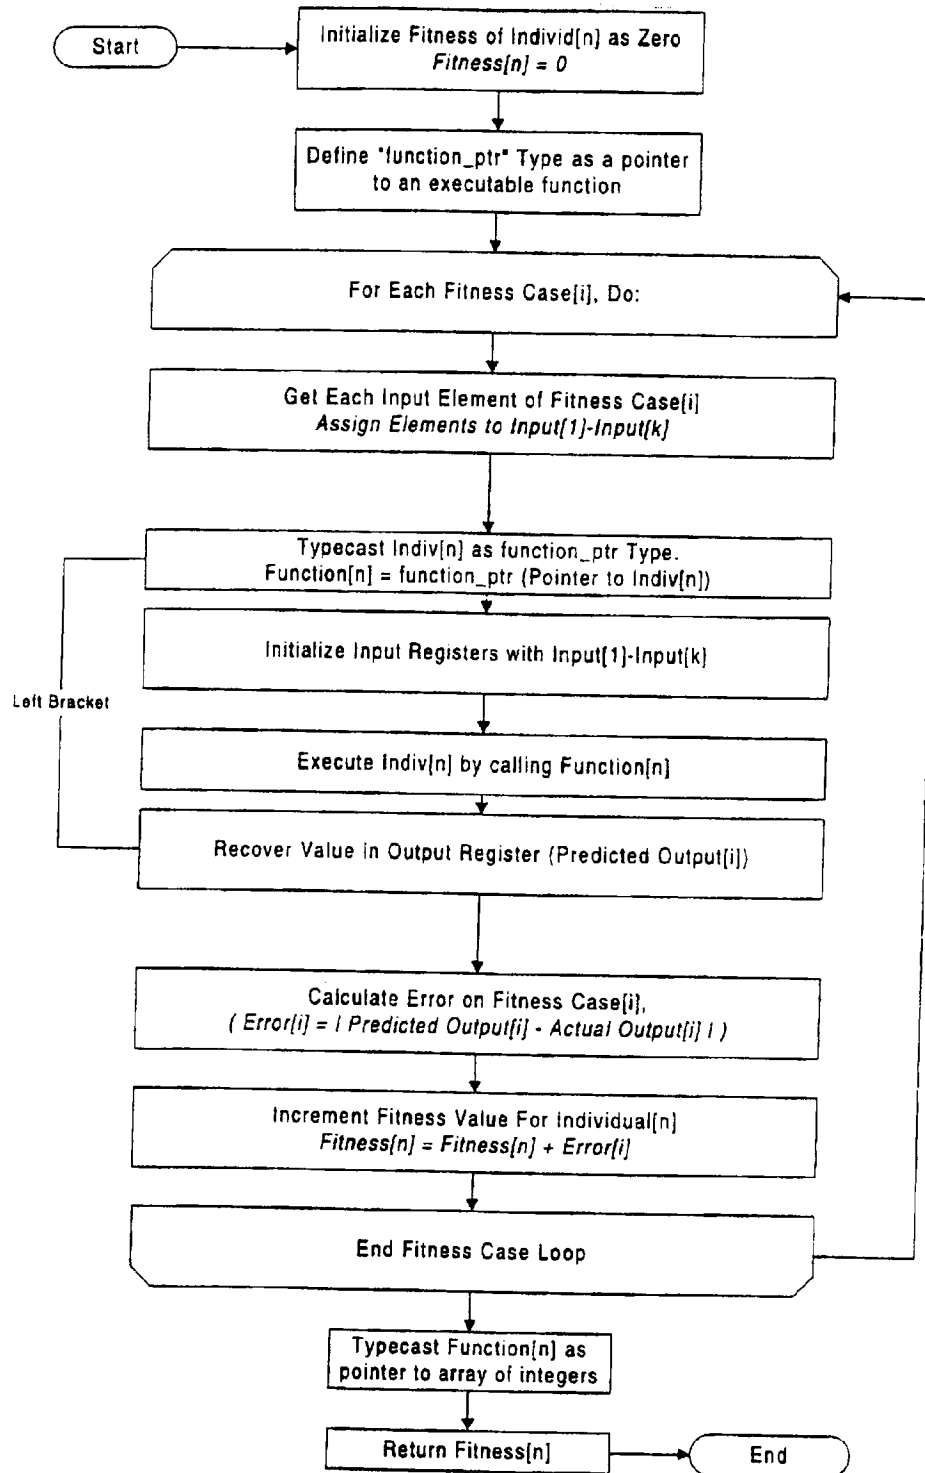


Figure 22g

39/46

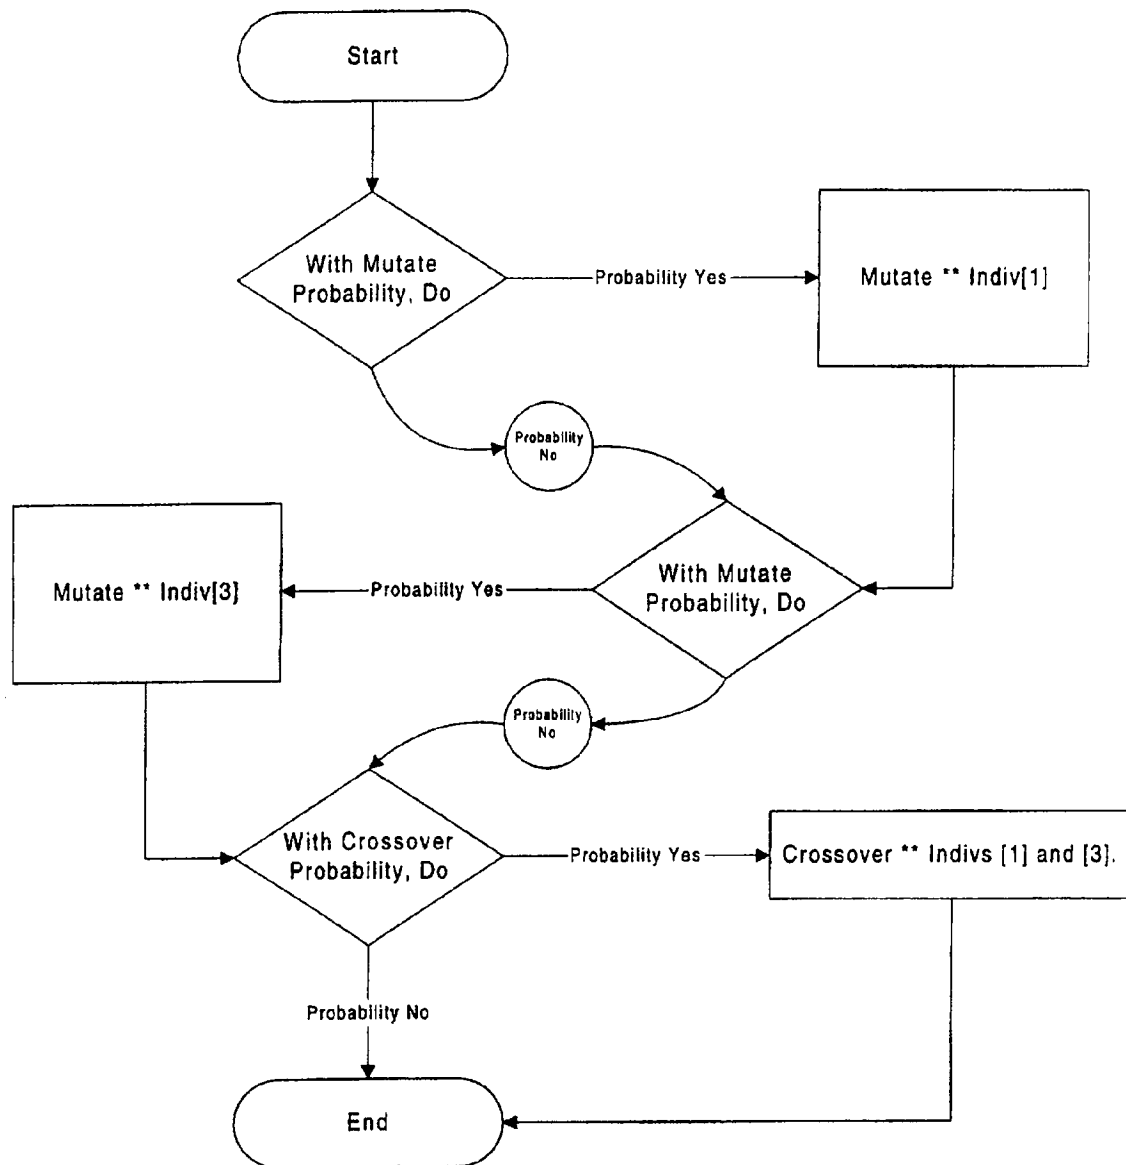


Figure 22h

40/46

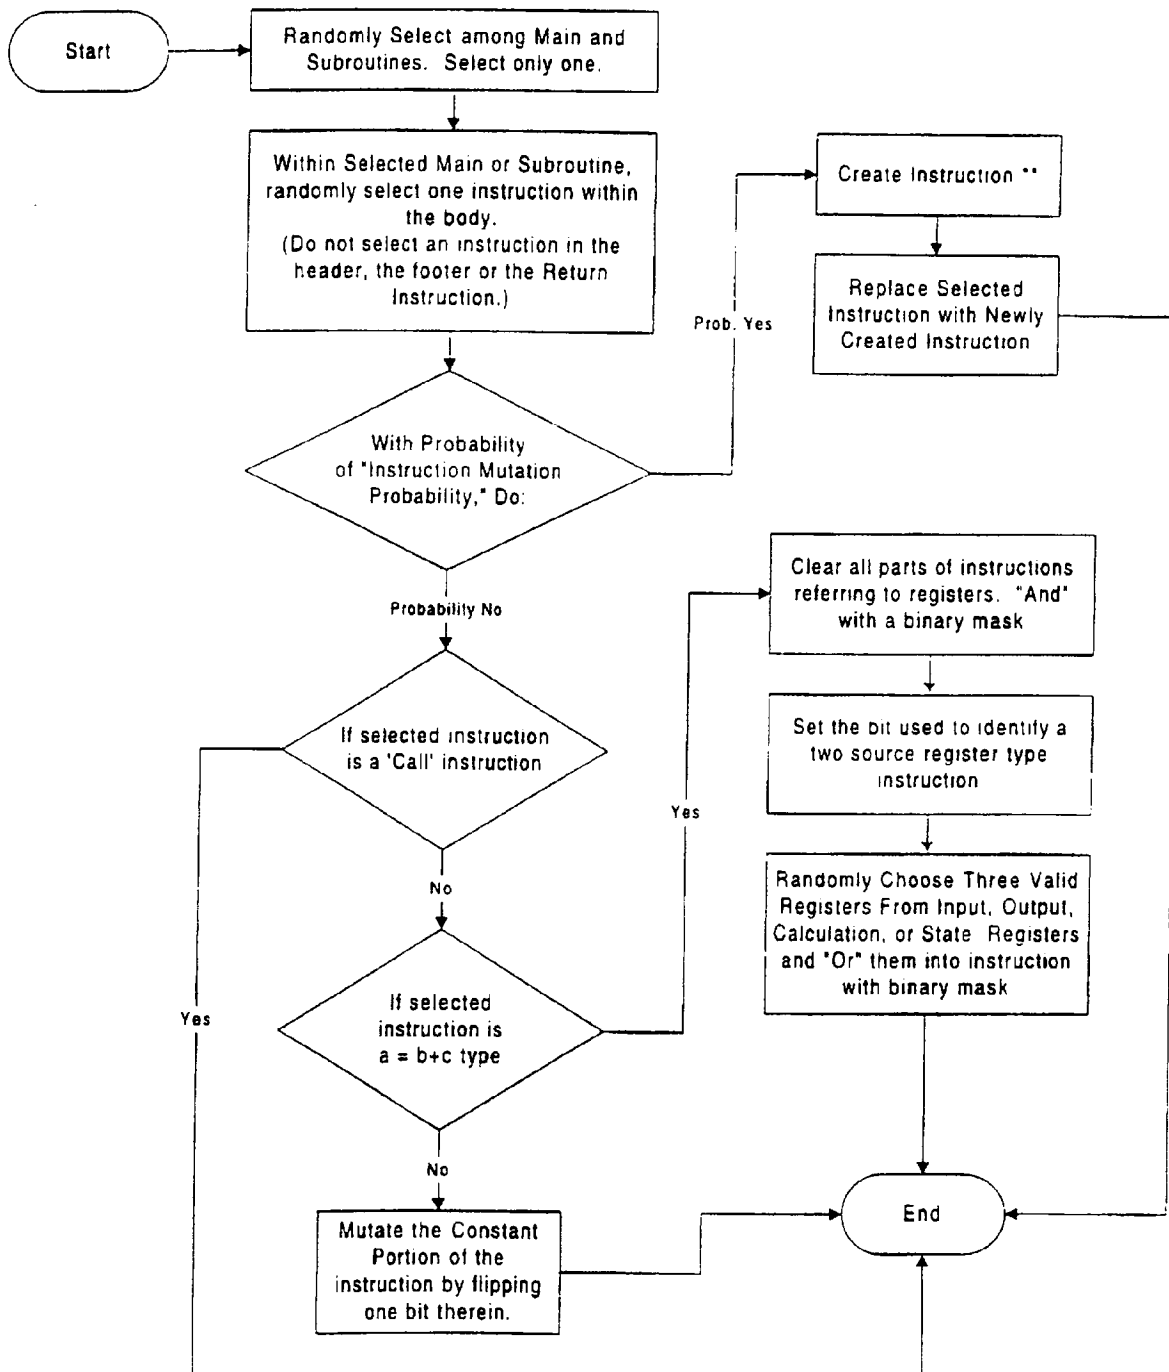


Figure 22i

41/46

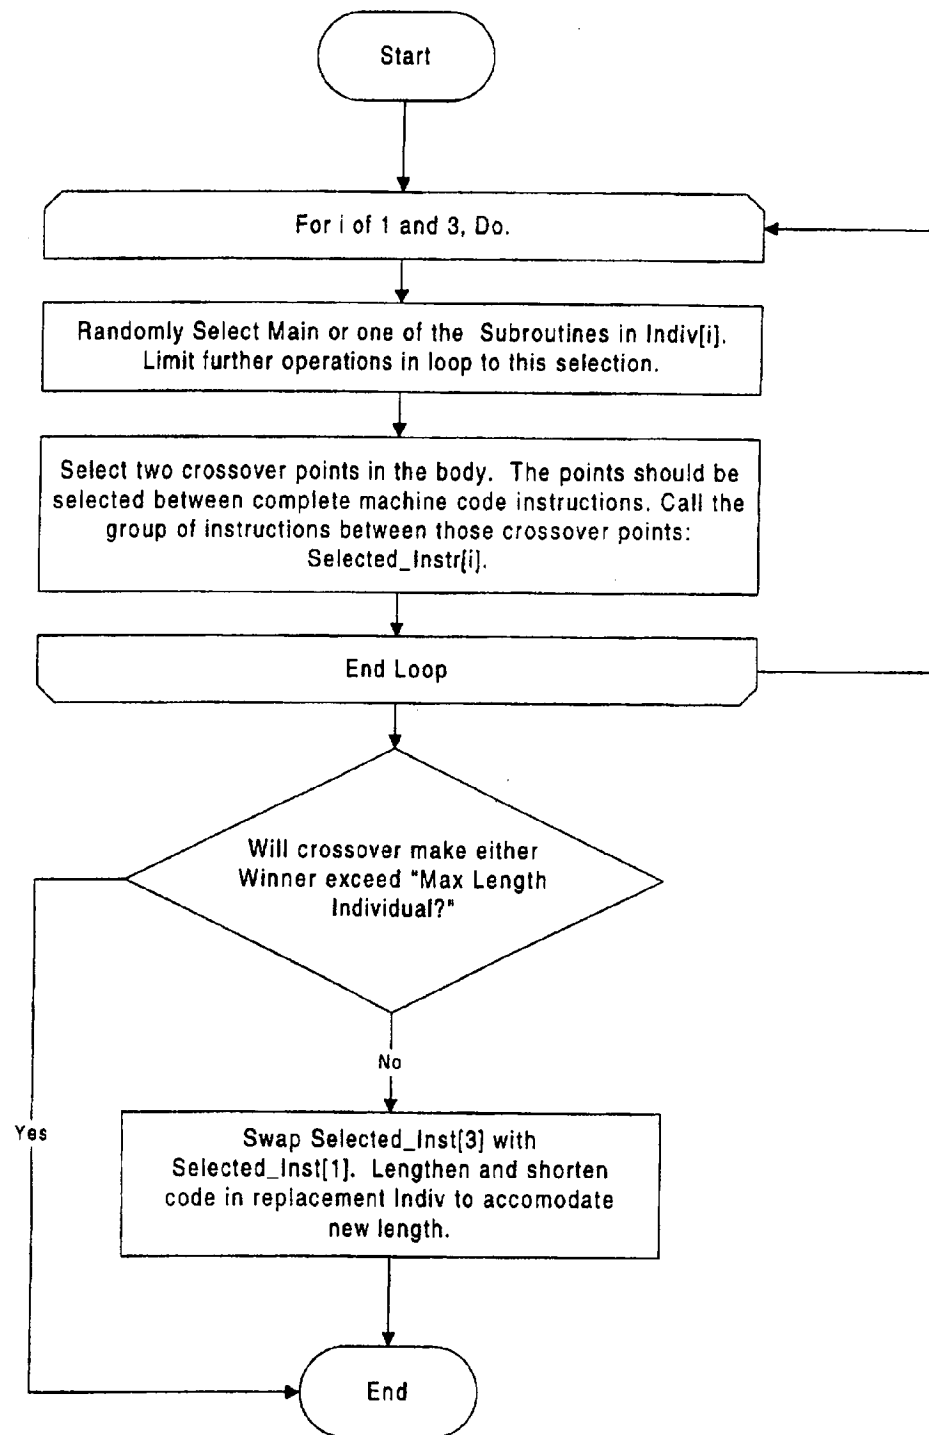


Figure 22j

42/46

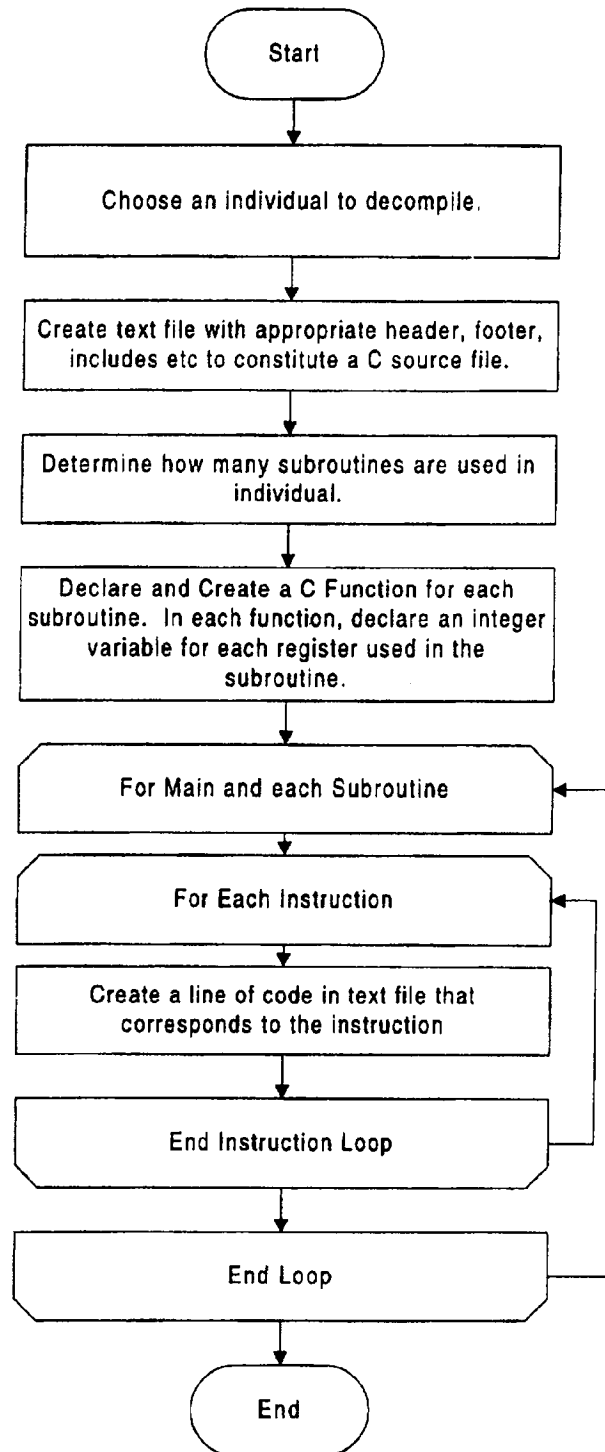


Figure 22k

43/46

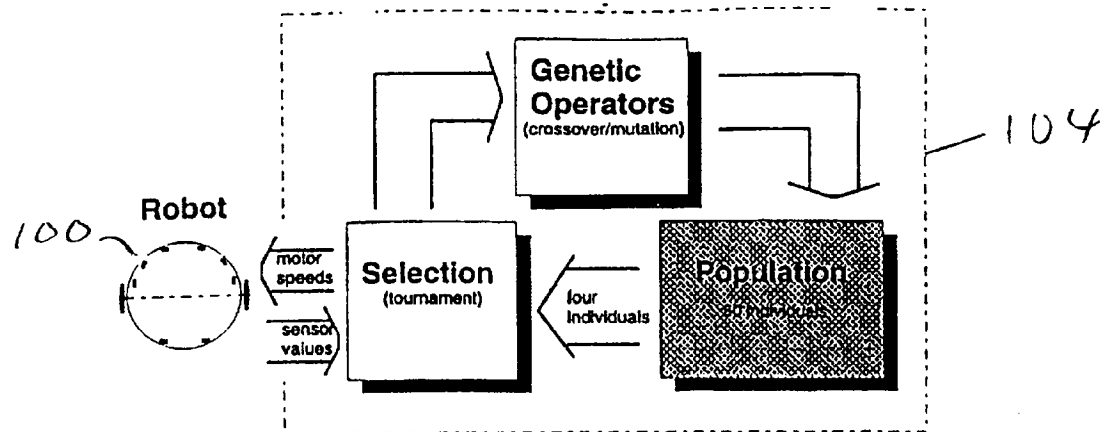
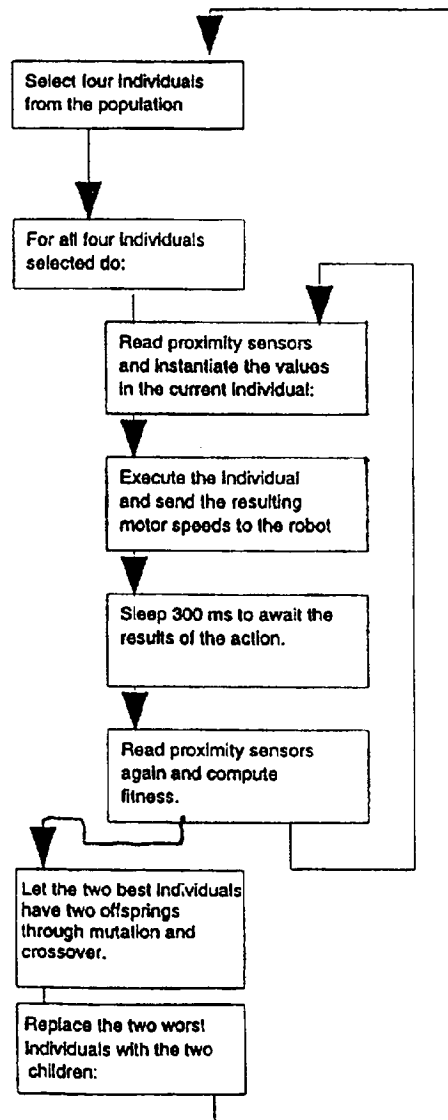


Fig. 23

Fig. 24
44/46




Fig. 25

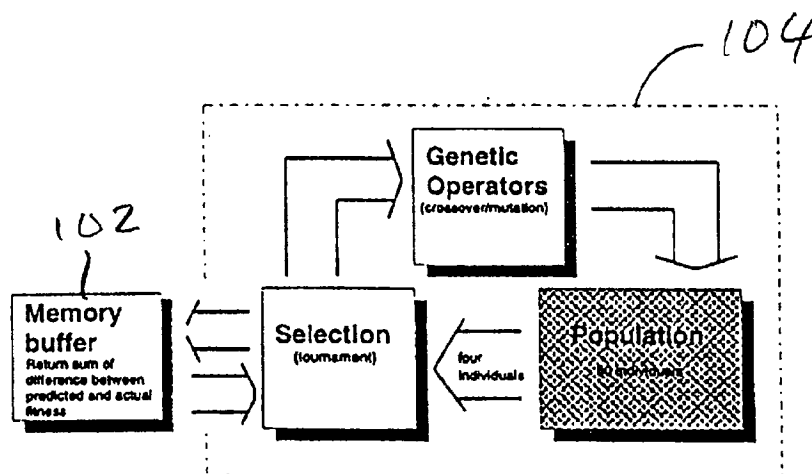


Fig. 26

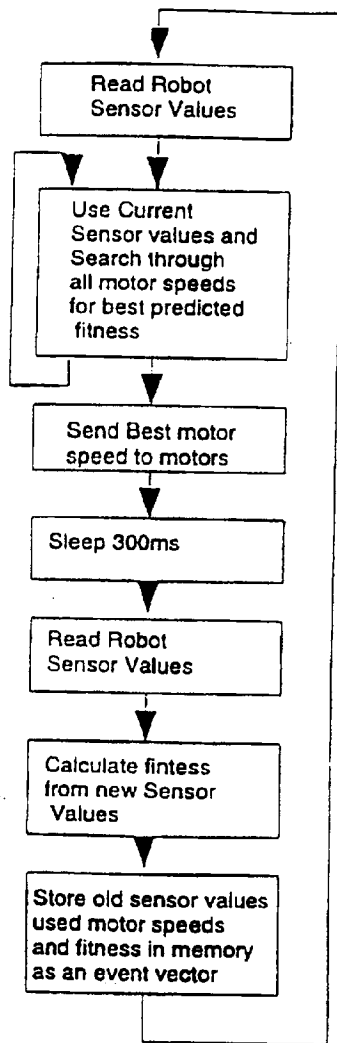


Fig. 27

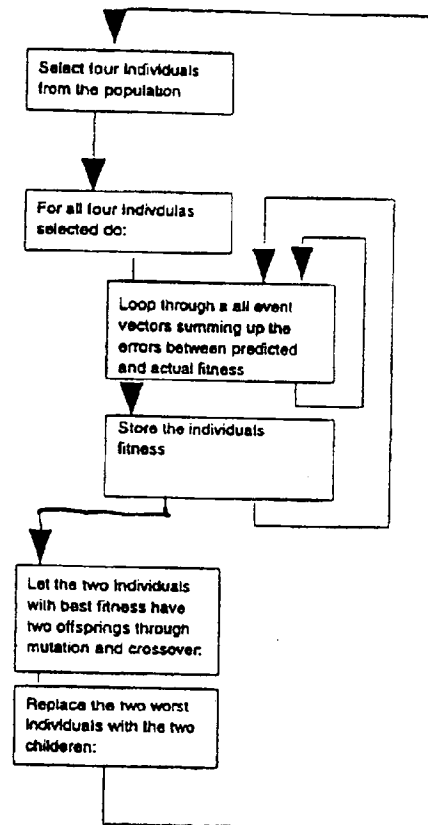


Fig. 28

46/46